Probes of new physics at the intensity and energy frontiers

Dedicated to my mother Vivianne Magill
and my father Walter Magill

Probes of new physics at the intensity and energy frontiers

by

Gabriel Magill, M.Sc.

A Thesis
Submitted to the School of Graduate Studies
in Partial Fulfillment of the Requirements
for the degree Doctor of Philosophy

McMaster University

Doctor of Philosophy, Department of Physics and Astronomy, McMaster University (2018), Hamilton, Ontario, Canada.

Title: Probes of new physics at the intensity and energy frontiers

Author: Gabriel Magill, M. Sc.
(University of Waterloo & Perimeter Institute for Theoretical Physics)

supervisors: Dr. Cliff Burgess & Dr. Itay Yavin

number of pages: 145

Lay Abstract

In this thesis, we review the basics of particle physics at neutrino experiments and particle colliders. We then motivate and develop key searches that can be performed to look for new particles at a series of existing and future experiments. We focus on new particles with masses between the electron mass and 1000 times the proton mass. The many searches we consider involve looking for processes that produce a single ray of light, a single electron, a pair of oppositely charged 'electron-like' events, new collider signatures, and/or modifications of star explosions. In the process, we set novel new bounds on many theories. We also motivate the construction of a new detector at the Large Hadron Collider called milliQan, and perform a full simulation to assess its future performance. The milliQan experiment has since been approved and is currently undergoing construction.

Abstract

In this thesis, we review the basics of phenomenology in particle physics at neutrino beam dump experiments and hadron colliders. We then consider the phenomenology of various new particles, with masses between 1 MeV and 1 TeV, at the intensity and energy frontiers. We perform sensitivity analyses for physics beyond the standard model at particle colliders (LEP and LHC) and a number of past and future neutrino beam dump experiments (SHiP, DUNE, LSND, MicroBooNE, MiniBooNE and SBND). In particular, we motivate searches for new heavy neutral leptons in single photon events at neutrino and collider experiments (and also via supernova cooling), millicharged particles in single electron events at neutrino experiments, lepton flavor violating scalars via standard model induced mixed flavor neutrino trident production at neutrino experiments, and colored scalar doublets at colliders in events with many jets, soft leptons and low missing energy. In the process, we set novel new bounds on the parameters of these theories and propose powerful new searches that can be performed. We also motivate the construction of a new detector at the LHC called milliQan, and perform a full Geant4 simulation to calculate its projected sensitivity for millicharged particles. The milliQan experiment has since been approved and is currently undergoing construction.

Acknowledgments

First and foremost, I would like to thank my parents, Vivianne and Walter Magill. Starting from a young age and throughout my upbringing, they've instilled in me very important skills that have gotten me through the various challenges I've faced in my life. They taught me the importance of hard work, attention to detail, that if you set your mind to something you can achieve it, and most importantly, to never give up. Throughout my studies, they've dedicated countless hours to helping me with school work, applications and extra-curricular activities, not to mention their large financial investment in helping me pay my university degree. I would certainly not be who I am or where I am today if it wasn't for my parents. They have done the same for my sister Tara Magill, who much better than me reflects all of these values, and is a constant source of inspiration. So thank you.

I am also greatly indebted to my partner and fellow physicist Natacha Altamirano. She has always supported and accepted me, and encouraged me to do new things. It has been a joy to discuss with her all matters of life, to travel with her, to discover new hobbies and new countries. Her way of thinking and analyzing topics never ceases to amaze me.

I would also like to thank my close collaborators Ryan Plestid, Maxim Pospelov and Brian Shuve. Most of the skills, practices and tools relevant to research in particle physics that I have learned have been fostered through them. They have always been very patient and helpful in our discussions, and have always been willing to contribute a lot of work and energy into our projects. I would also like to thank my supervisors Itay Yavin and Cliff Burgess for many discussions and guidance. In addition to their invaluable mentoring, they made a great effort to expose me to various groups, collaborators, collaborations and conferences. This close integration with the particle physics research community has been invaluable in shaping me as a researcher.

A city is only as good as the people in it. My life in Waterloo would certainly not have been as interesting as it was without these people: Andrés, Lucía, Carolina, Dalimil, Julían, Marcelo, Leo, Belén, John, John, Ravi, Pablo, Peter, Markus, Teddy, Dan, Fiona, Ke, Gunjan, and Néstor. Thanks for the beers, road trips, asados, foosball, coffee, and good times.

I would like to thank Debbie, Rosemary, Tina and Cheryl for diligently taking care of things on the administrative side. Lastly, I would like to thank the Canadian taxpayers, who have supported me through scholarships and assistantships. I would not have been able to fund my Ph.D. if it weren't for the strong support Canada shows to the sciences.

Table of Contents

List of Figures

List of Tables

Acronyms

BBN Big Bang Nucleosynthesis.

BNB Booster Neutrino Beam.

BSM beyond the standard model.

CERN Conseil Européen pour la Recherche Nucléaire.

COM center of mass.

DIS deep inelastic scattering.

DUNE Deep Underground Neutrino Experiment.

DY Drell-Yan.

EPA equivalent photon approximation.

HNLs heavy neutral leptons.

IP interaction point.

LAr-TPC liquid argon time projection chamber.

LHC Large Hadron Collider.

LSND Liquid Scintillator Neutrino Detector.

mCPs millicharged particles.

NTP neutrino trident production.

PDF probability distribution function.

PMNS Pontecorvo-Maki-Nakagawa-Sakata.

PMTs photomultiplier tubes.

POT protons on target.

SBN Short Baseline Neutrino.

SBND Short Baseline Near Detector.

SHiP Search for Hidden Particles.

SM standard model.

SN1987A Supernova 1987A.

Declaration of Authorship

This is a sandwich thesis based on 6 papers that appear as follows:

- Chapter 2

  Austin Ball et al. "A Letter of Intent to Install a milli-charged Particle Detector at LHC P5," (2016), arXiv:1607.04669 [ins-det], CMS IN-2016/002

  Gabriel Magill. A Milli-Charged Particle Detector at LHC P5. DESY-PROC, pages 39-42, Hamburg, Jun 2017. 12th Patras Workshop on Axions, WIMPs and WISPs, Jeju Island (South Korea), Verlag Deutsches Elektronen-Synchrotron.

- Chapter 3

  Gabriel Magill, Ryan Plestid, Maxim Pospelov, and Yu-Dai Tsai, "Millicharged particles in neutrino experiments" (2018), arXiv:1806.03310 [hep-ph] (submitted for publication in Phys. Rev. Letters)

- Chapter 4

  Gabriel Magill, Ryan Plestid, Maxim Pospelov, and Yu-Dai Tsai, "Dipole portal to heavy neutral leptons," (2018), arXiv:1803.03262 [hep-ph] (submitted for publication in Phys. Rev. D)

- Chapter 5

  Gabriel Magill and Ryan Plestid, "Neutrino trident production at the intensity frontier," Phys. Rev. **D95**, 073004 (2017), arXiv:1612.05642 [hep-ph]

- Chapter 6

  Gabriel Magill and Ryan Plestid, "Probing new charged scalars with neutrino trident production," Phys.Rev. **D97** (2018), arXiv:1710.08431 [hep-ph]

- Chapter 7

  Gabriel Magill, Brian Shuve, Jonathan Kozaczuk, Eder Izaguirre, "Probing Compressed States at the LHC," (2018), in preparation

In the 1st paper, Gabriel Magill developed the Geant4 simulation (in collaboration with James London in the early stages), and the generation of the kinematic distributions of milli-charged particles at the LHC interaction point using MadGraph. He also wrote the section in the Letter of Intent pertaining to these sections. The passage of the milli-charged particles from the LHC interaction point to the detector was done by Bennett Marsh and Frank Golf. Obtaining a realistic fully functional Geant4 simulation of the milliQan detector was one of the biggest and most important components of the paper, which required over a year of work.

In the 2nd paper, Ryan Plestid was responsible for upgrading the meson decay packages from the 3rd paper. Gabriel Magill was responsible for developing and implementing the pipeline to perform the sensitivity studies at the intensity frontier experiments. Yu-Dai performed the Drell-Yan calculations. Ryan and Gabriel both contributed significantly to the writing of the paper.

In the 3rd paper, Ryan Plestid significantly developed the meson decay packages and optical depth calculations, and wrote the sections pertaining to supernovae as well as other sections of the paper. Maxim Pospelov entirely wrote the BBN section as well as parts of the introduction and conclusion. Gabriel Magill worked primarily on all of the other calculations that appear in the paper. These include doing all the sensitivity studies at LEP, LHC, MiniBooNE, MicroBooNE, SBND, SHiP, LSND, doing the coherent, diffractive and Drell-Yan calculations, coding the meson production phase space distributions (in association with Ryan) and writing the sections pertaining to these topics. Gabriel also contributed significantly to the emissivity and optical depth calculations of SN1987A. Yu-Dai worked on the emissivity calculations.

In the 4th and 5th papers, both authors contributed significantly to the calculations, writing and editing of the paper.

In the 7th paper, Gabriel Magill performed almost all of the signal and background simulations, and developed most of the analysis code for the recast and projected sensitivity study (with help from Brian Shuve). Gabriel also contributed significantly to the writing of the draft. Jonathan calculated the decay length and branching ratio for the new colored scalars, and simulated the $2j$ QCD background process.

Chapter 1

Introduction

1.1 The search for new physics

From the discovery of the X-Ray in 1895 by W.C. Röntgen [1] up until the discovery of the Higgs boson in 2012 at the Large Hadron Collider (LHC) [2, 3], we have built up our understanding of the particles that make up matter. These particles and their interactions are the basis of the standard model (SM) of particle physics. Its matter content as we currently understand it is summarized in Fig. 1.1.

The detection of the Higgs boson was the last missing piece of the SM (apart from neutrino masses), and widely established its validity at least as an effective theory of nature. Currently, searching for physics beyond the standard model (BSM) takes place on at least three main fronts:

- *Intensity frontier*: High luminosity and statistics experiments probe weakly interacting signals.
- *Energy frontier*: High energy collider experiments probe massive particles.
- *Cosmic frontier*: Earth based experiments study signals coming from space.

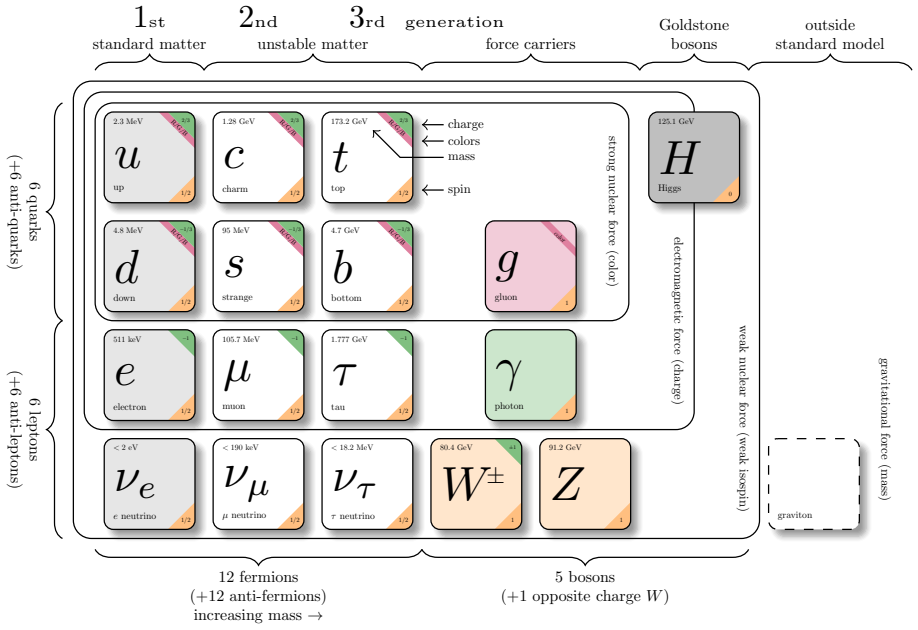

FIGURE 1.1: Particle content of the Standard Model that has been experimentally discovered. Credit to [4] for the figure.

Most of the focus of this thesis will be to motivate searches for new particles at the intensity and energy frontiers. There is currently a renewed interest in the intensity frontier related to neutrino measurements [5, 6, 7, 8, 9]. Indeed, neutrino masses are currently unaccounted for in the SM and require a mechanism to generate them. Studying flavor oscillations and neutrinos is extremely important as it allows one to determine mass differences between neutrinos, the neutrino mass hierarchy, theoretical models of mass generation, the charge parity violating structure of the leptonic sector, and more accurate measures of the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix [10, 11, 12]. The upcoming flagship experiment responsible for performing many of these measurements will be the Deep Underground Neutrino Experiment (DUNE) long-baseline experiment [5, 6]. By accelerating protons at 80 GeV on a graphite target, they will produce a flux of neutrinos that will first be detected at a 30T near detector in Fermilab (Illinois, USA), and later on at a 40kT far detector located 1300km away at the Sanford Underground Research Facility (South Dakota, USA). We note though that the DUNE collaboration is still considering different choices of detector technology and total mass.

Oscillation measurements have already revealed interesting hints that can potentially be interpreted as new physics. Two measurements that have received a lot of attention are from the MiniBooNE and Liquid Scintillator Neutrino Detector (LSND) collaborations [13, 14, 15], as potentially providing evidence for eV sterile neutrinos [16, 17] or heavy neutral leptons [18, 19]. At MiniBooNE, an excess of 3σ was observed for $\nu_\mu \to \nu_e$ appearances and at LSND, a significance of 3.8σ was observed for $\overline{\nu}_\mu \to \overline{\nu}_e$ appearances. The MiniBooNE result has recently been upgraded to a 4.8σ measurement [20]. The interpretation of these measurements is scrutinized [21] due to the lack of ability of the Cherenkov detectors being used by each collaboration to distinguish photons from electrons. Due to the importance of these measurements, Fermilab is currently building a Short Baseline Neutrino (SBN) program [9] that will consist of 3 experiments fed by neutrinos coming from the Booster Neutrino Beam (BNB) [22]. The SBN program consists of three experiments — Short Baseline Near Detector (SBND), MicroBooNE and ICARUS — which are respectively located in a line at distances of 110m, 470m, and 600m from the BNB proton target. MiniBooNE was an old experiment on the same line at approximately 540m from the target.

The experiments mentioned above, along with a number of others summarized in Table 1.1, are of broader interest than the main measurements for which they are designed. Indeed, many of these experiments have a substantial number of protons on target (POT), large detection volume and are built on excellent detector technology involving a liquid argon time projection chamber (LAr-TPC) or a nuclear emulsion detector. This opens up the possibility of using them to search for new physics that interacts with neutrinos, electrons or photons. We now give a brief overview of the theories that we will be interested in probing at the intensity frontier (and at the energy frontier for higher masses). For a more detailed discussion on how these theories are related to subsequent chapters, see Section 1.3.

We first consider millicharged particles (mCPs). mCPs χ are particles of arbitrary mass that have a coupling to photons of strength ϵe, where $\epsilon < 1$ and e is the electric charge carried by an electron. Strictly speaking, the term *milli* implies that $\epsilon \in [0.001, 0.01)$, although we will use this name regardless of the value of ϵ. They are of interest because finding them would imply that electric charge for free particles is not quantized, or more conservatively that the unit of quantization if it exists must be smaller or equal to ϵ [23, 24]. Furthermore,

they are of interest because the mechanism of kinetic mixing that can give rise to mCPs is general, natural, and contains very few assumptions [25].

We will consider searching for mCPs at the intensity frontier via meson decays to off-shell photons. When the masses of these particles exceeds half the meson mass and/or are above the center of mass (COM) energy of the neutrino experiment, mCPs can be produced at particle colliders. At the LHC, detection is very difficult since the small ionization energy deposits they produce get washed out by hadronic activity [26]. The collider bounds from CMS only exclude $|\epsilon| \geq 2/3$ for masses below 310 GeV, and $|\epsilon| \geq 1/3$ for masses below 140 GeV. Therefore we propose an experiment called milliQan [27, 28, 29]. With a prototype currently under construction, it aims to search for mCPs with $|\epsilon| \gtrsim 10^{-3}$ for masses of $\mathcal{O}(100)$ MeV, and $|\epsilon| \gtrsim 10^{-2}$ for masses of $\mathcal{O}(10)$ GeV. Elimination of the hadronic background will be achieved by locating the experiment in a tunnel at a distance of 33m from the LHC's interaction point (IP), 17m of which is through rock and cement.

Secondly, we will be interested in heavy neutral leptons (HNLs). HNLs [18, 19, 30, 31] can couple to both neutrinos and photons with a dimension-5 magnetic dipole moment. They can be a possible explanation of both the LSND and MiniBooNE excesses by acting as a source of single photons (which are misidentified as single electrons) in Cherenkov detectors. We will show that future LAr-TPC detectors at neutrino experiments will be very valuable in probing these particles in the single photon final state. If the HNLs' couplings to photons and neutrinos are sufficiently small and their masses are below ~ 100 MeV, they can provide an additional source of supernova cooling and modify the observed time signal of neutrinos arriving on earth [32, 33] compared to what was observed [34, 35]. We will use the Supernova 1987A (SN1987A) explosion to derive bounds [36, 37]. In a similar coupling and mass range, HNLs can modify the proton to neutron ratio during the Big Bang Nucleosynthesis (BBN) era [30], which we consider as well.

Lastly, we can consider new charged scalars coupling to neutrinos and charged leptons [38, 39, 40, 41, 42]. Charged scalars such as the Zee-Babu or Higgs triplet models are used as mechanisms to generate neutrino masses that don't rely on sterile neutrinos. Furthermore, their interaction pattern with leptons could shed light on the flavor structure of the SM. At neutrino experiments, they can be

Beam Dumps	$\mathbf{E_{POT}}$	$\mathbf{L_{Target}}$	**# POT**	**Technology**
LSND [49]	1.74 GeV	26m*	1.69×10^{23}	CH_2 Cherenkov
MiniBooNE [50]	9 GeV	540m	5.58×10^{20}	CH_2 Cherenkov
MicroBooNE* [51]	9 GeV	470m	$1.32 * 10^{21}$	LAr-TPC
SBND* [52]	9 GeV	110m	6.6×10^{20}	LAr-TPC
DUNE* Near [5]	80 GeV	600m	3×10^{22}	LAr-TPC
SHiP* ν_τ [8]	400 GeV	50m	2×10^{20}	Pb/Emulsion
Colliders	$\sqrt{s}$	$\mathbf{L_{IP}}$	$\mathcal{L}$ **[fb^{-1}]**	**Technology**
LHC [53, 54]	8(13) TeV	$\sim$ 5cm	19.7(36.1)	pp Collider
LEP [55, 56, 57]	91.2GeV	$\sim$ 5cm	0.1	e^+e^- Collider
milliQan* [28]	14 TeV	33m*	3000	Scintillator/PMT

TABLE 1.1: List of experiments considered in this work. Experiments with an asterisk are either planned or in construction, and lengths with an asterisk denote an experiment located off-axis. For future experiments, the data sample size is the expected value for the lifetime of the experiment. For existing experiments, the information shown corresponds to analyses we have recasted. The Search for Hidden Particles (SHiP) and collider experiments are located in Conseil Européen pour la Recherche Nucléaire (CERN), LSND is located in Los Alamos, and the remaining beam dump experiments are located in Fermilab. In addition to these experiments, we also consider bounds from cooling of supernova SN1987A, and the proton to neutron ratio during BBN.

detected via neutrino trident production (NTP) [43, 44, 45]. This process was previously observed [46, 47, 48] in the muon-only channel, with rates consistent with SM predictions. By calculating SM rates of mixed flavor NTP at future experiments, we can set projected limits on new charged scalars.

1.2 Phenomenology in particle physics

Most of the work appearing in the subsequent chapters of the thesis relate to searching for new particles at various experiments. Towards this, it will be instructive to perform a comprehensive overview of phenomenology in particle physics. We will first take an experiment independent approach and describe the concepts of statistical inference for establishing sensitivity. Next, we will discuss how to apply these concepts in particle physics, focusing on neutrino beam dump experiments and hadron colliders.

1.2.1 Statistical inference

We begin by discussing statistical inference as it relates to particle physics. Sensitivity analyses in the context of statistical inference have been fundamental for understanding the matter content of the SM, as well as the interactions between each of its particles. Given the importance of finding new particles towards establishing the full theory of nature, it is necessary to have a rigorous framework for establishing the presence of new particles. Our discussion hereafter will be a literature review based on [58, 59, 60, 17]. In particle physics, the criteria for discovery is the so-called 5σ condition. Given a hypothesis H_0 which contains the SM along with all of its particles and interactions as we currently understand them, this condition in short states that a discovery is made if the data deviates from the H_0 predictions by 5 standard deviations. The data in this case is usually a set of events $\boldsymbol{n}$ in a histogram. The prediction is how many events we would expect to see according to H_0. In defining an event, there are various measurements $\boldsymbol{x}$ that are made and a set of cuts $\mathcal{D}$, imposed on $\boldsymbol{x}$, chosen so as to maximize background rejection and signal acceptance. After running the experiment for a long time, we can combine all of the events (consistent with the cuts $\mathcal{D}$) into a test statistic t. This test statistic is a random variable, and is distributed according to a probability distribution function (PDF). The test statistic is often a χ^2 value or a likelihood function. The likelihood function $L(\boldsymbol{\theta}, \boldsymbol{\nu}) = P(\boldsymbol{n}|\boldsymbol{\theta}, \boldsymbol{\nu})$ is obtained by regarding the joint probability P as a function of only the parameters $\boldsymbol{\theta}$ and $\boldsymbol{\nu}$, where the observations $\boldsymbol{n}$ have been fixed by the experiment. The set of $\boldsymbol{\theta}$ are model parameters that need to be fixed or whose values we are testing (i.e. coupling constants, particle masses, etc) and which directly affect the predictions. These are the parameters that are immediately of interest and for which our measurement is geared to. The $\boldsymbol{\nu}$ are nuisance parameters that also need to be fixed but that are not of immediate interest, such as systematic uncertainties, calibrations, reconstructed energy profiles, luminosity uncertainties, and so on.

As an example, we consider a counting experiment with N observations/bins, each described by Poisson statistics. In this case, the likelihood function is given as

$$L(\boldsymbol{\theta}, \boldsymbol{\nu}) = \prod_{i=1}^{N} \frac{\mu_i^{n_i}}{n_i!} e^{-\mu_i}, \tag{1}$$

where the μ_i are average expected event counts for each bin i, and depend on $\boldsymbol{\theta}, \boldsymbol{\nu}$. These averages are typically obtained by doing a large ensemble of simulations using the H_0 model.

At this point, we briefly divert to make an interesting observation. In cases where we are interested in inclusive measurements ($N = 1$), Eq. (1) reduces to the typical Poisson distribution with an average prediction of μ and a standard deviation of $\sqrt{\mu}$. The observed data n is therefore $(n-\mu)/\sqrt{\mu}$ standard deviations away from what is expected according to H_0. If it is 5 standard deviations away, the measurement is considered significant enough to be a discovery. The number of standard deviations $(n-\mu)/\sqrt{\mu}$ is an estimate of the significance of the measurement, can be rewritten as the commonly used $S/\sqrt{B}$, where S (B) refers to the signal (background) counts[1]. It is clear that when the backgrounds are small, this metric for calculating signal sensitivity is pathological. For large backgrounds, a direct comparison to the more formal procedure described in [58, 30] show that the $S/\sqrt{B}$ approximation gives an accurate significance to within 10% for backgrounds of 100. When $N > 1$ and for large backgrounds, the sensitivity expression reduces to a χ^2 test. This is known as Wilks' theorem [61].

Going back to Eq. (1), the dependence on $\boldsymbol{\nu}$ can be removed by the following procedure. Suppose we can study the effects of $\boldsymbol{\nu}$ by making measurements $\boldsymbol{y}$ in a control region that is independent of $\boldsymbol{x}$, then we can write

$$L(\boldsymbol{\theta}, \boldsymbol{\nu}) = P_x(\boldsymbol{x}|\boldsymbol{\theta}) P_y(\boldsymbol{y}|\boldsymbol{\nu}). \tag{2}$$

For each value of $\boldsymbol{\theta}$, we can find the $\boldsymbol{\nu}$ which maximizes L, and we are left with a likelihood function $L(\boldsymbol{\theta})$. Recalling that the likelihood function was our test statistic t randomly distributed according to a PDF, and assuming that larger values of t show disagreement with our model, we can test the compatibility of its measured value t_{obs} with the model hypothesis H_0 using the p-value

$$\text{p-val} = \int_{t_{\text{obs}}}^{\infty} \text{PDF}(t|H_0) dt. \tag{3}$$

Then, if p-val is smaller than some pre-determined value, the H_0 hypothesis is rejected. For the 5σ test, this corresponds to a p-value of 2.87×10^{-7}.

[1] This is true when the errors are statistics dominated. When they are systematics dominated and we don't expect the sensitivity to improve with more statistics, the sensitivity becomes S/B.

We now discuss how constraints for new physics are obtained. Instead of the SM hypothesis H_0, we consider the hypothesis H_1 which contains the SM augmented with new physics. The likelihood function under H_1 can be written as

$$L(\boldsymbol{\theta}) = \frac{(s+b)^n}{n!} e^{-(s+b)}, \tag{4}$$

or its appropriate generalization for multiple bins. The averages are now the expected background b from the the SM and the additional signal rate s from new physics. Introducing priors $\pi(s)$ (or $\pi(\boldsymbol{\theta})$) on the new physics, and an overall normalization X, we can use Bayes' theorem and introduce a posterior probability

$$p(\boldsymbol{\theta}|\boldsymbol{n}) = \frac{L(\boldsymbol{\theta})\pi(\boldsymbol{\theta})}{X}. \tag{5}$$

We can then find the smallest new physics signal s_{up} compatible with the observations and expectations at $1-\alpha$ credibility level, by solving

$$\alpha = \int_{s_{\text{up}}}^{\infty} p(\theta|\boldsymbol{n}) ds. \tag{6}$$

Finally, we can directly relate s_{up} to the parameters of the new physics to set constraints on them. We now describe how this is done.

1.2.2 Neutrino beam dumps

In the previous section, we have provided an overview of how to exclude or prove the existence of new physics. This involved introducing measurements $\boldsymbol{x}$, a set of cuts $\mathcal{D}$ on these measurements and a signal rate s_{up} consistent with a given observation and expectation. Now, we will describe more concretely the meaning of these quantities in the context of neutrino experiments, and the main considerations that go into determining them. In the next section, we will do the same for collider experiments.

All of the neutrino experiments we consider in this work are based on a proton beam dump and proceed as follows:

1. Protons are accelerated towards a target to produce mesons and Drell-Yan collisions, among other things. These are referred to as POT.

2. An assortment of prompt, long-lived, charged, and neutral mesons are produced. Some experiments, such as DUNE and the SBN facility, have a decay chamber and a magnetic field following the proton target, which allows to focus long-lived charged mesons such as $\pi^{\pm}$ and $K^{\pm}$. These are the source of the experiments' (anti-)neutrino flux. Other experiments such as SHiP have a beam stomp which effectively allows only prompt mesons such as η and π^0 to decay.

3. Neutrinos (and any other weakly interacting particles) that are produced travel through dirt or shielding material until they arrive to the detector. Along the way, neutrinos can oscillate. The probability of this happening [10] is largely controlled by $\Delta m^2_{ij} L/4E$, namely the mass separation between two neutrinos, the distance through which they travel and their energy.

4. Once at the detector, they interact with either the electrons, partons, nucleons or nuclei in the detector. This interaction can consist of either a single energy deposit at one location, or several energy deposits with a given spatial and temporal separation (this is often the case for processes involving muon production).

The collection of all of these steps go into determining the expected backgrounds and signals rates, and performing the experiment yields the observation. The relevant formula summarizing each of these steps can approximately[2] be expressed as

$$\text{Rates} = \sum_{\text{Energies}} N_X(E_i) \times \frac{N_{\text{Target}}}{\text{Area}} \times \sigma^X_{\text{Target}}(E_i) \times \mathcal{E}. \tag{7}$$

In Eq. (7), $N_X(E_i)$ represents the number of particles of type 'X' with energy E_i arriving at the detector, $\sigma^X_{\text{Target}}(E_i)$ is the detection cross section of particle X with a target in the detector (i.e. an electron, nucleon, nuclei, etc) consistent with the angular and recoil cuts in the experiment, N_{Target} is the total number of such targets inside the active volume of the detector and $\mathcal{E}$ is a detection efficiency. The "Area" in (7) stands for the active volume divided by the average length $\langle l \rangle$ traversed by particles inside the detector. The particle type X depends

[2]This is true when the detector is sufficiently far away from the meson source and the physics interacting in the detector is sufficiently weakly interacting. In general, $N_X(E_i)$ should also be a function of angle and we should include an exponential function encoding the mean free path of X with the detector.

on the process we are considering. For oscillation experiments, X will be (anti-)neutrinos. For HNLs produced by neutrinos interacting inside the detector, X will also be a neutrino. For HNLs produced directly from mesons and decaying inside the detector, X will be the HNLs. Relating this to the previous section, by setting the Rates in Eq. (7) equal to s_{up}, we can set bounds on the parameters contained in $\sigma^{X}_{\text{Target}}(E)$.

We now provide additional details on the items listed above. We begin the discussion at the detector level, and will work our way backwards to the POT. We define an *event* as the set of all *objects* detected according to the requirements of $\mathcal{D}$, such that the number of observations made increases by one. The objects seen in the detector involve SM leptonic activity, hadronic activity, or a combination of both. The presence or absence of hadronic activity will depend largely on the momentum exchange q of the incoming particle with the nucleus. For concreteness, a depiction of an incoming neutrino interacting via a photon mediator to produce HNLs is shown in Fig. 1.2.

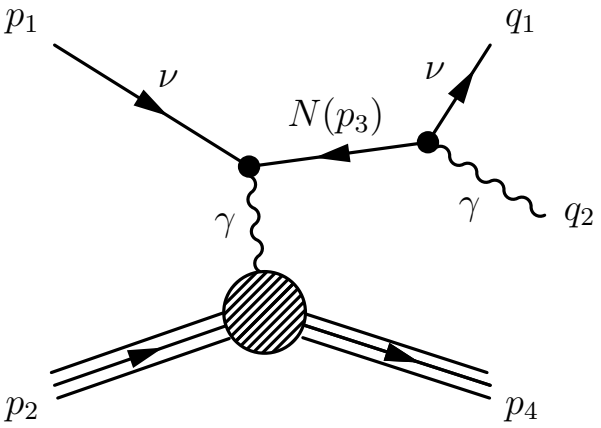

FIGURE 1.2: A neutrino scattering with matter to produce a heavy neutral lepton, which decays after travelling some distance to a photon and a neutrino. Depending on the magnitude of the momentum transfer in the photon, the reaction can proceed coherently with the nucleus, diffractively with nucleons or in the deep inelastic regime with the partons.

When the magnitude of the momentum transfer $Q \equiv |q|$ is below $\Lambda_{\text{QCD}}/(A)^{\frac{1}{3}}$ (where Λ is the QCD scale and A is the number of nucleons), the photon cannot resolve individual nucleons and the reaction proceeds coherently over the whole nucleus. This leads to a Z^2 enhancement with the number of protons. The form factor of the nucleus $F(Q^2)$ enforces this range for Q. We use the Woods-Saxon

form factor

$$F_{\rm WS}(Q^2) = \frac{1}{N}\mathcal{F}\left\{\frac{V_0}{1+\exp\left(\frac{r-r_0A^{1/3}}{a}\right)}\right\} \tag{8}$$

with $\mathcal{F}$ denoting the Fourier transform with respect to r, and N is a normalization factor given by the volume integral over the nuclear charge distribution [62]. The various parameters are set as $r_0 \approx 1.126$ fm, $a \approx 0.523$ fm, and $V_0 = (4\pi A r_0^3/3)^{-1}$. Different choices of form factor modify the result on the 10% level. Coherent scattering typically leads to little hadronic activity, since the kinetic recoil energy of a non-relativistic nucleus scales quadratically with Q and inversely with its mass.

For intermediate momentum transfers Q in between $0.217/(A)^{\frac{1}{3}}$ GeV and M_p, we have diffractive scattering and the photon scatters with individual nucleons. In these cases, the nucleons that are ejected lead to observable signatures and can constitute one of the objects of the event. The cross section scales linearly in the number of protons and neutrons, and the momentum profile can be represented by the form factors F_1 and F_2. They take on different values depending if they are for the neutron or proton. Their values are given [63, 64] by solving the system of equations

$$\begin{aligned} G_{\gamma,E}^{p,n} &= F_{1,\gamma}^{p,n} - \frac{Q^2}{4M_{\rm Nucleon}^2}F_{2,\gamma}^{p,n} \\ G_{\gamma,M}^{p,n} &= F_{1,\gamma}^{p,n} + F_{2,\gamma}^{p,n} \end{aligned} \tag{9}$$

with

$$\begin{aligned} G_{\gamma,E}^{\{p,n\}} &= \{G_D, 0\} \\ G_{\gamma,M}^{\{p,n\}} &= \mu_{\{p,n\}}G_D \\ G_D &= \frac{1}{(1+Q^2/0.71\ {\rm GeV}^2)^2} \\ \mu_{p,n} &= \{2.793, -1.913\}. \end{aligned} \tag{10}$$

We then obtain

$$\sigma_{total} = Z \times \sigma_p + (A-Z)\times\sigma_n. \tag{11}$$

For $Q \gtrsim 1$ GeV $\approx M_p$, the scattering occurs with respect to the quarks and one must employ the deep inelastic scattering (DIS) or Drell-Yan (DY) frameworks.

The form factors one must use are the parton momentum distribution functions of quarks and gluons inside the nucleons [65]. These produce large deposits of hadronic activity. There are also other kinematic regimes which have not discussed, for instance resonant meson production. Lastly, we mention that there is overlap in Q between where one of these regimes end and the other begins.

Having given some examples of the the kinds of objects one can see, we must now specify the relevant $\boldsymbol{x}$ variables based on these objects, as well as any cuts or vetos $\mathcal{D}$ we impose on them. In practice, the variables that are available — such as the energy/momentum and the type of particle that has been reconstructed — depend on the type of detector being used. There are several detector technologies we will be considering in this work. We briefly describe each of them based on [66].

The first class of experiments are Scintillation+Cherenkov detectors (e.g. LSND) and Cherenkov detectors (e.g. MiniBooNE) [67]. The Cherenkov light inside a medium of refractive index n is produced by particles with speeds β larger than $1/n$. These produce cones of light with an opening angle (from the direction of the initial particle) of $\cos\theta = 1/n\beta$. The scintillation light is produced by energy deposits inside a material according to the characteristic scintillation response of the material. The light that is produced needs to be collected and detected, which is often done using photomultiplier tubes (PMTs). Cherenkov detectors unfortunately cannot distinguish an electron from a photon which undergoes conversion to an electron-positron pair, since both lead to the production of electromagnetic showers.

The second class of experiments are LAr-TPC (e.g. DUNE, SBND, MicroBooNE). These improve upon Cherenkov detectors in that they can resolve photons from electrons. They operate according to the principle by which particles produce ionization inside a counting gas or liquid. The ionization is then collected by wires using an $E-$ field and readout. LAr-TPC are one of the best technologies for track reconstruction.

Lastly, we will be considering nuclear emulsion technology for detectors (e.g. SHiP ν_τ detector). Nuclear emulsions are composed of small crystals inside a gelatine substrate. A particle travelling through this substrate alters the crystals and produces a track, which can then be seen in a lab offline. In the case of SHiP, the emulsion region is composed of a series of thin films that are stacked

into a group. Each group is interleaved with lead plates to induce a reaction. Nuclear emulsion detectors also provide excellent track reconstruction and vertex resolution.

Ideally, the purpose of these detectors is to measure energy deposits in order to fully reconstruct the spectrum (or flux) of the incoming particle. In practice, there are several intrinsic limitations which we briefly note that prevent one from accurately determining the spectrum. Firstly, some of the outgoing particles aren't detected or are lost, which leads to difficulties in reconstructing the energy of the incident particle. Secondly, a given angle and energy of a final state particle can map to subspace of angles and energies for the incoming particle. There are also non-trivial mappings, which are difficult to extract, from the true energy of the incoming particle to its reconstructed energy. To make matters worst, the energy spectrum of the incoming neutrinos is usually broad. To address some of these challenges, various techniques and observables have been developed [68]. Lastly, it is a currently a field of active research [69] how to simultaneously and efficiently model the momentum distribution of the nucleons inside the nucleus, inter-nucleon interactions, fermi blocking effects, and final state interactions within the nucleus. These naturally all lead to difficulties in applying conservation laws.

Fortunately, in phenomenological studies the flux of incident (anti-)neutrinos or new weakly interacting particles can be simulated (instead of reversed engineered). To do so, we must have analytic doubly differential cross sections $d^2\sigma/d\theta dp$ for the mesons outgoing momentum and angular profile. We then convolve this with the meson's partial decay width. We can then calculate an acceptance for the decay products' reaching the detector, and their energy spectrum. There are several different $d^2\sigma/d\theta dp$ fits for meson production, depending on the meson species and the proton beam energy. In Table 1.2, we tabulate all of the differential cross sections that are used in this work.

The momentum and angular dependence of these distributions is quite complicated. In order to gain some intuition, we show in Fig. 1.3 density plots for the BMPT, Burman-Smith and Sanford-Wang distributions, for π^0 production. These distributions are plotted respectively for the POT beam energy and material composition of the SHiP, LSND, and Fermilab SBN experiments. We focus on neutral pions since they are copiously produced and are often the dominant

Name	Mesons	Beam Energies	Relevant Experiments
Burman-Smith [70]	π	≤ 800 MeV	LSND
Sanford-Wang [22, 71]	π	$7 - 35$ GeV	Fermilab BNB
Feynman Scaling [72]	η, $K^{\pm}$	$8.89 - 24$ GeV	Fermilab BNB
BMPT [73]	π, η	$100 - 450$ GeV	SHiP, DUNE
GJK [74]	J/Ψ	$200 - 450$ GeV	SHiP, DUNE*

TABLE 1.2: List of selected mesons produced and their energy+angular spectrum distributions at a variety of beam dump experiments. We assume that the GJK distribution is still approximately correct at DUNE energies after rescaling the total production cross section as a function of s. The proton on target beam energies shown correspond to the data that went into determining the empirical fits. The mesons listed is not exhaustive and represents the mesons we have considered for our processes. For instance, the BMPT distribution also describes $K^{\pm}$ production, however their decays aren't prompt enough to contribute significantly at the SHiP experiment. Kaons can also be described using Sanford-Wang distributions.

source of production for new particles that can be pair-produced from an off-shell photon via Dalitz-like decays. There are two features we can notice from these plots. Firstly is the steep cutoff of the density as a function of momentum. In fact, all of these distributions are exponentially decaying to some power of the momentum, in addition to various power law corrections. Secondly, there is the presence of a $\sin\theta$ measure, which removes the support of the differential distributions in the neighbourhood of $\theta = 0$.

Comparing now differences between the various distributions, we notice that at the low POT beam energies of LSND, π^0 production has substantial support at large angles from the beam line. This is important, because the LSND experiment is off-axis from the beam line by approximately 0.2 radians. It follows from the low energies of the protons impinging on the target that the pions produced will be fairly non-relativistic. At the other extreme, we have the BMPT distribution at SHiP. The pions that are produced have momentums mostly between $5 - 10$ GeV, and the distribution is peaked much more forward. This is important since the ν_τ detector at the SHiP experiment is located on axis at roughly twice the distance from the proton target as compared to LSND. Although the π^0 momentums at SHiP don't extend nearly as closely to E_{POT} as they did in the case of LSND, they will nevertheless be very relativistic. This implies that any decay products of π^0 at SHiP will also mostly be collimated in the forward direction.

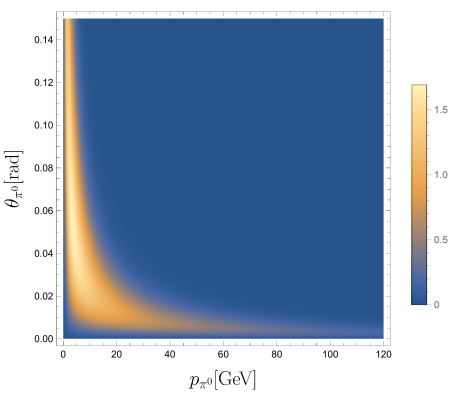

(a) BMPT $E_{\mathrm{POT}} = 400$ GeV

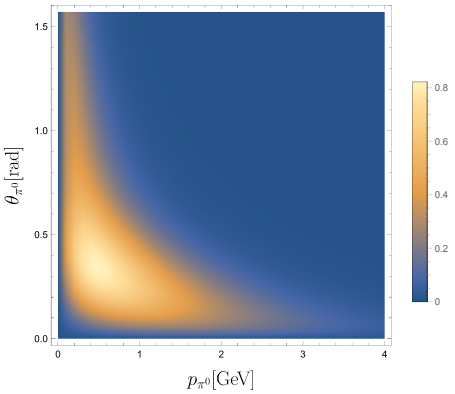

(b) Sanford-Wang $E_{\mathrm{POT}} = 9$ GeV

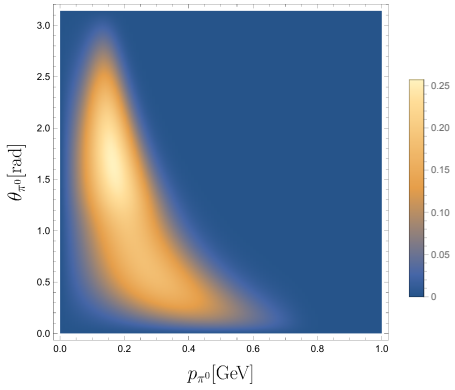

(c) Burman-Smith $E_{\mathrm{POT}} = 1.74$ GeV

FIGURE 1.3: Double differential distribution $d^2\sigma/d\theta dp$ for π^0 production normalized per proton on target. The BMPT distribution assumes SHiP POT energies and is calculated for a target composed of 50% Tungsten and 50% Molybdenum. The Sanford-Wang distribution assumes BNB POT energies and is calculated for a Beryllium target. The Burman-Smith distribution assumes LSND energies and is calculated for a target composed 50% of Water and 50% of Tungsten.

1.2.3 Particle colliders

At particle colliders, detection and exclusions also proceed according to the discussion in Section 1.2.1. The objects that make up an event however differ greatly from those in neutrino experiments. Both the ATLAS and CMS experiment at the LHC are cylindrical detectors that are made up of several layers [54, 53], each layer being optimized to detect a given type of particle and furnish a measurement in the set of $\boldsymbol{x}$.

The innermost layers provide tracking information for charged particles that is used to infer their momentum. Following the tracker is the electromagnetic calorimeter which is used to absorb electrons and photons and measure their energy. Then comes the hadronic calorimeter which does the same for hadrons. Lastly, there are muon spectrometers which are used to tag muons. There is also a solenoid magnet to generate a b-field, whose location differs between ATLAS and CMS. The LHC can therefore tag particles (with a given efficiency and misidentification rate) and determine whether it's a photon, electron, muon, or colored particle, and measure its energy, momentum, polar and azimuthal angle. We will not go into details about how this is done. Instead, we will highlight several LHC observables that are worth understanding for the phenomenological studies in this work.

Transverse Momentum: Although the momentum component along the beamline of particles can be measured, it is not often useful. The LHC is a proton-proton collider. However, in hard collisions, it is not the protons that collide, but rather the partons inside. The probability of a given parton participating in this collision and the fraction of the proton momentum that it carries is given by a parton distribution function [65]. Given the probabilistic nature of the quarks' and gluon's momentum along the beamline, it is the transverse momentum $p_T = \sqrt{p_x^2 + p_y^2}$ of objects in an event that is usually quoted. In contrast, for lepton colliders such as LEP, the initial momentum of each colliding beam is known.

Missing energy: Transverse to the beamline, momentum is conserved. That said, neutrinos and any other weakly interacting particles can escape detection, leading to an apparent momentum imbalance. The missing transverse momentum is defined [75, 76] as the negative vector sum of the transverse momenta of all

identified physics objects. The magnitude of the missing transverse momentum is known as the missing energy of an event.

Jets: In the hadronic calorimeters, quarks are not directly observed. Due to the confining nature of QCD, the quarks that are initially produced at the parton level after a collision hadronize to form bound states. During this process, they also shower additional QCD radiation [77]. This process leads to an (ideally) localized cluster of energy deposited in the hadronic calorimeter. Based on each energy deposit i in a given granularity resolution, jets are reconstructed according to a clustering algorithm. A jet can therefore be very crudely thought of as the energetic parent particle (or particles) that created a certain cluster of energy through showering and hadronization. The most popular clustering algorithm in use is the anti-k_T algorithm [78]. It proceeds as follows. Given an entity i in the hadronic calorimeter, the set of all other entities $\{k\}$ and the beam B, one defines in a particular way two metrics. From these, we calculate the distance d_{ij}, where $j \in \{k\}$ is the entity closest to i under the metric, and an 'absolute' distance d_{iB}. If $d_{ij} < d_{iB}$, the clustering algorithm combines the 4-momentums of i and j into a new entity. If $d_{ij} \geq d_{iB}$, i is declared as a jet and removed from the set $\{k\}$.

Simulations: For phenomenological studies, it is of interest to numerically study processes at the LHC. The simulation of LHC processes takes place over several steps.

1. A sample of events at parton level are generated for a given model and final state. A common software in use to do this is MadGraph [79]. Theoretical models are implemented using the Feynrules package and are loaded into MadGraph [80, 81].

2. Given an ensemble of events at parton level, one simulates their showering and hadronization. This is commonly done using Pythia [77], Herwig++ [82], and SHERPA [83]. In doing so, one must take care to include matching for jets. Matching is a Monte Carlo procedure necessary because identical QCD final states in certain regions of phase space can be generated both at the parton level and by showering. This leads to a double counting of the weight of that particular event.

3. Lastly, one must include detector effects. There are two ways to do this. One way is to parametrize detection efficiencies according to angles, momentum and particle ID. Secondly, one can directly simulate the detector response to the passage of particles through it. This is what we have done for the milliQan experiment using the Geant4 toolkit [84].

1.3 Outline of the thesis

Up until now, we have given a pedagogical overview of the background material that is needed to understand subsequent chapters. This included an overview of the BSM physics that we will be interested in probing, the experiments at which we will motivate these searches, the objects and variables $\boldsymbol{x}$ that are relevant for these searches at neutrino and collider experiments, and lastly an overview of the statistical framework we are using to characterize the sensitivity of these searches. We now describe the outline for the subsequent chapters.

In Chapter 2, we discuss the milliQan experiment to detect mCPs at the LHC, for which a prototype is currently being installed. At the time of writing, milliQan was a proposed experiment for which we constructed a fully functional Geant4 simulation, as well as a C++ analysis that reads in the ROOT files [85] generated in the simulation and performs a sensitivity study. This detector simulation was important to model certain effects, such as the attenuation length of photons in scintillators, deflections in the path of mCPs due to energy losses, and tests of various coincidence detection strategies. From a politics perspective, it acted as a proof of concept for our Letter of Intent, which was ultimately successful in gaining approval.

In Chapter 3, we discuss searching for mCPs at the intensity frontier. Surprisingly, there is not much literature in the way of using beam dump experiments as a production factory for mCPs. The results from the dedicated SLAC experiment [86] has set stringent bounds for masses and couplings relevant for beam dumps. That said, we show that we can already leverage existing data at LSND [87] and recast it to obtain bounds that slightly improve upon the SLAC limits in certain areas. Any improvements one can obtain is important, and very difficult since the bounds on ϵ scale roughly as the 8^{th} root of luminosity. Furthermore, we show how future neutrino experiments can outperform the SLAC

and LSND experiments, and provide complementary sensitivity to milliQan for masses between 0.1 GeV $-$ 5 GeV.

In Chapter 4, we motivate a search for heavy neutral leptons. We provide the first comprehensive sensitivity analysis for this model for masses between 1 MeV $-$ 10 TeV, obtaining limits from existing and projected experiments across colliders, neutrino experiments, SN1987A cooling and BBN. In the process, we rule out certain regions of interest for this model that were posited to explain the LSND and MiniBooNE anomalies.

In Chapter 5, we extend existing work on neutrino trident production. Standard model neutrino trident production was the object of several searches in the 90s [46, 47, 48], and was recently revived as a very powerful probe of Z' models gauged under $L_\mu - L_\tau$ [43]. Neutrino trident production involves the creation of a pair of oppositely charged leptons from an incoming neutrino. In previous work, the lepton pairs that were considered were only muon flavors. We extended this work by calculating production rates at SHiP and DUNE for mixed flavor pairs involving combinations of electrons, muons and tauons. We demonstrated that certain standard model mixed flavor trident modes can be produced at higher rates than conventional NTP.

In Chapter 6, we study how effective mixed trident production is in searching for new leptophilic charged scalars. We find that depending on the flavor structure of these new scalars, mixed flavor trident can easily outperform conventional NTP. In some cases, it also becomes very competitive with certain precision measurement bounds.

In Chapter 7, we present preliminary results of a proposed search for electrically charged scalars, which are also charged under $\mathrm{SU}(3) \times \mathrm{SU}(2)$. We motivate this search at the LHC, and consider a final state topology that has yet to be considered in existing searches. We show that this new topology, based on soft leptons, low missing energy, and many jets, can beat existing searches that we have recasted.

Lastly, in Chapter 8, we conclude and provide an outlook on the work presented in this thesis.

Chapter 2

Letter of intent to install a milli-charged particle detector at LHC P5

2.1 Preface

We propose a dedicated experiment to search for mCPs. The proposed experiment is a model-independent probe of mCPs, since it relies only on the masses and charges of such particles. A prototypical model of mCPs [88] consists of an abelian gauge field A'_μ that couples to a massive Dirac fermion ψ of mass $M_{\rm mCP}$ and to hypercharge via kinetic mixing. A simple field redefinition removes the kinetic mixing, yielding the Lagrangian

$$\mathcal{L} = \mathcal{L}_{\rm SM} - \frac{1}{4} A'_{\mu\nu} A'^{\mu\nu} + i\bar{\psi}\left(\not{\partial} + ie'\not{A}' - i\kappa e'\not{B} + iM_{\rm mCP}\right)\psi. \tag{12}$$

The new matter field ψ therefore has an electric charge of $\epsilon \equiv \kappa e' \cos\theta_W / e \ll 1$, a milli-charge. The parameter space spanned by the mass and charge of the mCPs is constrained by indirect observations [89, 90], although these can often be evaded by adding extra degrees of freedom. Previous experiments have also looked for non-quantized charged particles [26, 91]. However, the parameter space for mCPs with masses 0.1 GeV $\lesssim M_{\rm mCP} \lesssim$ 100 GeV and charges at the $10^{-3} \leq \epsilon \leq 10^{-1}$ level are largely unexplored by direct searches. In particular,

the CMS/ATLAS sensitivity to mCPs in monojet and missing energy searches is precluded by the penalty on the cross-section associated with initial state radiation. In the following, we summarize a recent proposal to construct a dedicated detector at the LHC to target this unexplored part of parameter space.

In order to evaluate the projected sensitivity of the experiment for various electric charges and masses, we performed a full simulation of the experiment, including a Geant4 [84] model of the detector. In the first stage, Madgraph5 and Madonia were used to simulate the production of mCPs via Drell-Yan, J/Ψ, $\Upsilon(1S)$, $\Upsilon(2S)$, and $\Upsilon(3S)$ channels at 14 TeV COM energy [79]. Particles produced at the interaction point were propagated using a map of the CMS magnetic field to the proposed experimental site. The effects of multiple scattering and energy loss were included. In the second stage, we calculated the signal efficiency by running the kinematic distributions of mCPs at the proposed experimental site through a full Geant4 simulation of the detector based on the specifications provided by the manufacturers for Saint-Gobain BC-408 plastic scintillators and Hamamatsu R329-02 PMTs [92, 93]. This was important to model the reflectivity, the light attenuation length, the dependence on small electric charge, and to test different coincidence strategies.

This paper appears on arXiv as

- Austin Ball et al. "A Letter of Intent to Install a milli-charged Particle Detector at LHC P5," (2016), arXiv:1607.04669 [ins-det]

and in an internal CMS note as

- CMS IN-2016/002.

A compressed version of this paper also appears in a conference proceeding as

- Gabriel Magill. A Milli-Charged Particle Detector at LHC P5. DESY-PROC, pages 39-42, Hamburg, Jun 2017. 12th Patras Workshop on Axions, WIMPs and WISPs, Jeju Island (South Korea), 20 Jun 2016 - 24 Jun 2016, Verlag Deutsches Elektronen-Synchrotron.

Contribution of author (reproduced from Declaration of Authorship): Gabriel Magill developed the Geant4 simulation (in collaboration with James London

in the early stages), and the generation of the kinematic distributions of milli-charged particles at the LHC interaction point using MadGraph. He also wrote the section in the Letter of Intent pertaining to these sections. The passage of the milli-charged particles from the LHC interaction point to the detector was done by Bennett Marsh and Frank Golf. Obtaining a realistic fully functional Geant4 simulation of the milliQan detector was one of the biggest and most important components of these papers, which required over a year of work.

2.2 Paper

A Letter of Intent to Install a Milli-charged Particle Detector at LHC P5

Austin Ball,[1] Jim Brooke,[2] Claudio Campagnari,[3] Albert De Roeck,[1] Brian Francis,[4] Martin Gastal,[1] Frank Golf,[3] Joel Goldstein,[2] Andy Haas,[5] Christopher S. Hill,[4] Eder Izaguirre,[6] Benjamin Kaplan,[5] Gabriel Magill,[7,6] Bennett Marsh,[3] David Miller,[8] Theo Prins,[1] Harry Shakeshaft,[1] David Stuart,[3] Max Swiatlowski,[8] and Itay Yavin[7,6]

[1] *CERN*

[2] *University of Bristol*

[3] *University of California, Santa Barbara*

[4] *The Ohio State University*

[5] *New York University*

[6] *Perimeter Institute for Theoretical Physics*

[7] *McMaster University*

[8] *University of Chicago*

(Dated: July 19, 2016)

Abstract

In this LOI we propose a dedicated experiment that would detect "milli-charged" particles produced by pp collisions at LHC Point 5. The experiment would be installed during LS2 in the vestigial drainage gallery above UXC and would not interfere with CMS operations. With 300 fb^{-1} of integrated luminosity, sensitivity to a particle with charge $\mathcal{O}(10^{-3})$ e can be achieved for masses of $\mathcal{O}(1)$ GeV, and charge $\mathcal{O}(10^{-2})$ e for masses of $\mathcal{O}(10)$ GeV, greatly extending the parameter space explored for particles with small charge and masses above 100 MeV.

I. INTRODUCTION

In an effort to explain galactic dark matter, the idea of additional "sectors" beyond the SM has received considerable attention with a variety of examples (hidden valleys, secluded sectors, dark sectors, hidden sectors). The experimental searches for evidence of these sectors are driven largely by two factors: the particular way by which the extra sector is coupled to the SM; and the mass scale(s) in the extra sector. The specific details of the extra sector (the precise gauge group, the number of extra particles, etc.) are often not of great importance in designing the search.

In this LOI we propose a dedicated experiment to search for milli-charged particles (mCP). Since the search depends only on the mass and charge of such particles, one can view this search as part of the general program to search for additional sectors but to give a concrete example, we consider an extra abelian gauge field that couples to a massive Dirac fermion ("dark QED") and that mixes with hypercharge through the kinetic term [1],

$$\mathcal{L} = \mathcal{L}_{\rm SM} \; + \; \mathcal{L}_{\rm extra-sector} \tag{1}$$

$$\begin{aligned} \mathcal{L}_{\rm extra-sector} \; = \; & -\frac{1}{4} A'_{\mu\nu} A'^{\mu\nu} + i\bar{\psi}\left(\partial\!\!\!/ + ie' A\!\!\!/' + iM_{\rm mCP}\right)\psi \\ & -\frac{\kappa}{2} A'_{\mu\nu} B^{\mu\nu}. \end{aligned} \tag{2}$$

Here ψ is a Dirac particle of mass $M_{\rm mCP}$ that is charged under the new $U(1)$ field A'_μ with charge e', and the field-strength is defined as $A'_{\mu\nu} = \partial_\mu A'_\nu - \partial_\nu A'_\mu$. The last term in Eq. (2) is a kinetic mixing term between the field strength of the new gauge boson and that of hypercharge [21].

Eliminating the mixing term by redefining the new gauge boson as, $A'_\mu \rightarrow A'_\mu + \kappa B_\mu$ results in a coupling of the charged matter field ψ to hypercharge (as well as an immaterial redefinition of the hypercharge coupling),

$$\begin{aligned} \mathcal{L} \; = \; & \mathcal{L}_{\rm SM} - \frac{1}{4} A'_{\mu\nu} A'^{\mu\nu} \\ & + \; i\bar{\psi}\left(\partial\!\!\!/ + ie' A\!\!\!/' - i\kappa e' B\!\!\!/ + iM_{\rm mCP}\right)\psi. \end{aligned} \tag{3}$$

The new matter field ψ therefore acts as a field charged under hypercharge with a charge $\kappa e'$, a milli-charge [1]. The mCP ψ couples to the photon and Z^0 boson with a charge

$\kappa e' \cos\theta_W$ and $-\kappa e' \sin\theta_W$, respectively. The fractional charge in units of the electric charge is therefore $\epsilon \equiv \kappa e' \cos\theta_W / e$, where $\epsilon \ll 1$.

Previous experiments have looked for non-quantized charged particles [2–5]. The parameter space spanned by the mass and charge of the mCPs is also constrained by indirect observations from astrophysical systems [3, 6–8], the cosmic microwave background [9, 10], big-bang nucleosynthesis [11], and universe over-closure bounds [6]. While direct searches robustly constrain the parameter space of mCPs, indirect observations can be evaded by adding extra degrees of freedom. In particular, the parameter space for mCPs with masses $M_{\rm mCP}$ $0.1 \lesssim M_{\rm mCP} \lesssim 100$ GeV is largely unexplored by direct searches.

A natural question is whether the current general-purpose LHC experiments can improve sensitivity in the parameter space for mCPs with masses $M_{\rm mCP}$ $0.1 \lesssim M_{\rm mCP} \lesssim 100$ GeV. A recent analysis looking for low ionizing particles in CMS excluded particles with charge $\pm e/3$ for $M_{\rm mCP} < 140$ GeV and particles with charge $\pm 2e/3$ for $M_{\rm mCP} < 310$ GeV [5]. For fractional charges much smaller than unity, mCPs could be detected in monojet plus missing energy searches. However, for electroweak production of new quasi-invisible states, such as the mCPs studied here, the penalty on the cross-section associated with initial state radiation precludes CMS/ATLAS sensitivity to such particles, even with the very large datasets envisaged for the HL-LHC. Thus, to detect mCPs at the LHC, an alternative experimental strategy is needed.

In Ref. [12] some of the authors of this LOI proposed a new search to be conducted at the LHC with a dedicated detector targeting this unexplored part of parameter space, namely mCP masses $0.1 \lesssim M_{\rm mCP} \lesssim 100$ GeV, for charges Q at the $10^{-3}\ e - 10^{-1}\ e$ level. The experimental apparatus envisaged in this paper was one or more scintillator detector layers of roughly 1 m^3 each, positioned near one of the high-luminosity interaction points of the LHC. The experimental signature would consist of a few photo-electrons (PE) arising from the small ionization produced by the mCPs that travel unimpeded through material after escaping the ATLAS or CMS detectors. The proposed experiment is a model-independent probe of mCPs, since it relies only on the production and detection of mCPs through their QED interactions.

We have found the PX56 Observation and Drainage gallery above the CMS underground experimental cavern (UXC) to be an ideal site for such an experiment, and we propose in the following to situate there our experimental apparatus, which we call milliQan.

II. SITE SELECTION

LHC Point 5, Cessy is host to the Compact Muon Solenoid (CMS) experiment and its supporting infrastructure. As such, the appropriate services are available for the installation and operation of milliQan. We propose that the detector be set in the PX56 drainage gallery located above CMS UXC. Material access is limited by a door measuring 0.9 m $\times$ 2.1 m. That door links the drainage gallery and the platform installed in the shaft. This platform is within the coverage area of the overhead crane installed in the assembly hall building (SX5). Any components larger than a regular toolbox will have to be lowered to the gallery using the overhead crane and then passed through this door. Personnel access will be through a sector door located between the PM54 staircase and the drainage gallery. The door is interlocked with the LHC access control system. This entails that no access to milliQan will be possible while the LHC is in operation. The electrical infrastructure available in the CMS service cavern (USC55) can be expanded to bring power to milliQan using existing cable channels in the PM54 shaft. Our studies indicate that forced air from a portable air conditioning unit should provide sufficient cooling to hold the PMTs somewhat below room temperature and reduce the backgrounds from dark noise.

The proposed gallery is limited in space. Knowing this, a 3D model was combined with a laser scan of the gallery to give a best as-built estimation. Using the model an optimal position was found which offered; a distance to the IP of 33 m, 17 m of which is through rock, an angle of 43.1 degrees from the horizontal plane. Under these conditions clearance between the corners of the detector and the gallery can be as little as 30 mm. The selected location in the 3D model is shown in Figure 1. Clearance increases as the detector is moved away from the IP so the final position may come down to a compromise. These values are current best estimates; some small amounts of variation in the final design should be expected.

While the installation of the detector is not anticipated to present major challenges, a detailed work sequence is being studied. It will allow identifying the logistics operations required to move the components to their final destination in the drainage gallery and the corresponding necessary tooling. The study will also investigate a workflow for assembly, mounting and then alignment of the detector. A preliminary assessment shows that the CERN survey group will be able to align milliQan with a precision of about 2 cm. Safety

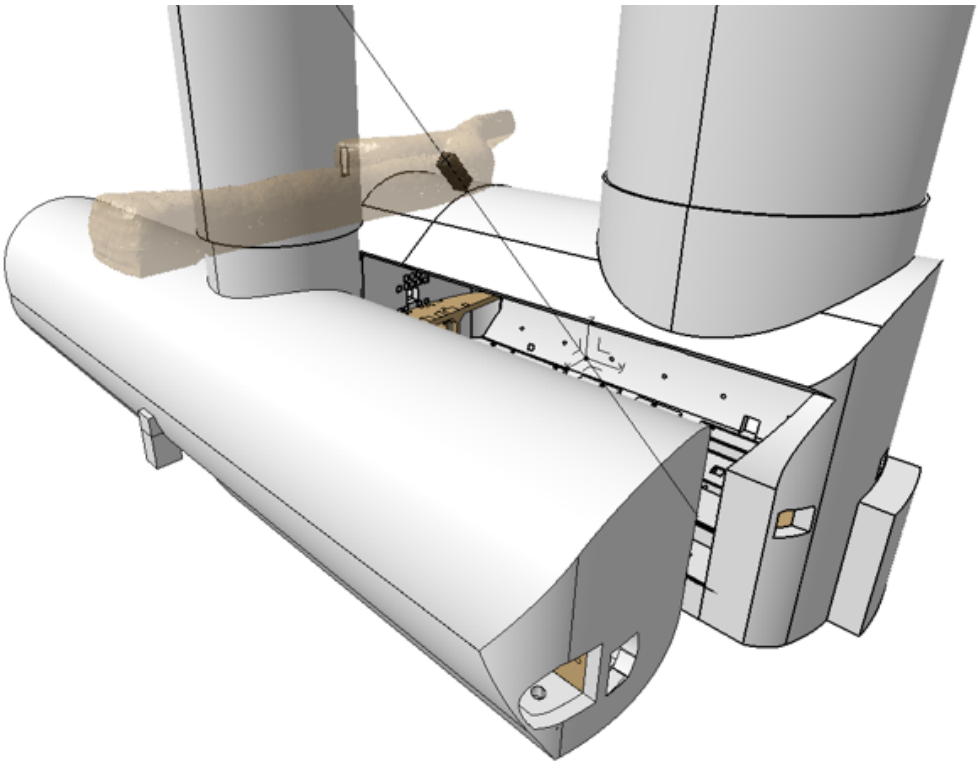

FIG. 1: 3D model showing optimal position of milliQan within the PX56 Drainage and Observation gallery located above CMS UXC.

aspects associated with installation, operation and maintenance of the detector will also have to be further studied. Fire loads and access requirements are all constraints to be considered. It is already certain that the detector support structure must allow a passage of at least 60 cm through the gallery. During access periods, this will require the detector to be moved into a "stowed" position where the length of the detector will be parallel with the wall of the gallery.

In terms of beneficial occupancy of the drainage gallery, this proposal has already been run through various LHC operation bodies and no conflicts were identified. In particular, we have received confirmation from the LHC integration office that they have no plans to ever install anything in the drainage gallery and that from their point of view it is available for use for milliQan. While still a lot of design and integration work has to be done, it seems that this scheme has a strong chance of meeting the requirements to install milliQan at point 5 in a reasonably short time.

III. RELATIONSHIP WITH CMS

A design requirement is that milliQan would be operated with no interference to CMS operations. The milliQan detector will self-trigger to a dedicated readout as described in Section VII. The only critical needs from CMS would be basic infrastructure such as power and ethernet, and read access to the CMS BRIL DB in order to receive information on the luminosity delivered to Point 5. The experiment would also need access to the LHC clock, which could be provided by a TCDS fiber. TCDS fibers can also provide run/luminosity section/orbit markers that would be used to synchronize milliQan with the CMS luminosity system. Finally, while not a requirement, access to the BPTX signals via LEMO cables would be useful to be able to distinguish collisions from empty buckets.

IV. DETECTOR CONCEPT

The detector that we propose to install at the location described in Section II is a 1 m $\times$ 1 m $\times$ 3 m plastic scintillator array. The array will be oriented such that the long axis points at the nominal CMS IP. The array is subdivided into 3 sections each containing 400 5 cm $\times$ 5 cm $\times$ 80 cm scintillator bars optically coupled to high-gain PMTs. A triple-incidence within a 15 ns time window along longitudinally contiguous bars in each of the 3 sections will be required in order to reduce the dark-current noise, which we expect to be the dominant background [22].

A $Q = 1e$ minimum-ionizing charged particle leaves roughly 2 MeV/cm in a material of density 1 g/cm^3 [13]. For plastic scintillator, such energy deposition results in about 10^4 photons per MeV, meaning 1.6×10^6 photons would be liberated in a 80 cm long scintillator. For a particle with electric charge $Q < 1e$, the energy deposition is reduced by the factor of Q^2 resulting, for fractional charges at the lower limit of the sensitive range, in just a few photons liberated on average in the same 80 cm long scintillator. Allowing for an estimated detection efficiency of about 10%, we therefore expect an average of $\mathcal{O}(1)$ photoelectron (PE) from each attached phototube for each mCP with $Q = \mathcal{O}(10^{-3})$ e that traverses our 80 cm plastic scintillator [2]. The signal we will search for is a longitudinal triple-incidence of hits with one or more PEs. Requiring triple-incidence will control the background to $\mathcal{O}(10)$ events per year with $N_{\mathrm{PE}} \geq 1$, as discussed in Section VIII.

V. MECHANICS, COOLING, AND MAGNETIC SHIELDING

We are developing the mechanics to support the detector in the drainage gallery site so as to allow both modular assembly and movement of the detector to a stowed position during access periods.

We have a working design consisting of three stacks, one for each of the three layers along the mCP flight direction, that can be separately assembled. The scintillator bars are mounted on trays within each stack as sketched in Fig. 2. The three stacks, each of which might weigh up to $\sim$ 2000 kg, will be mounted on an adjustable platform that can tilt the full assembly to point toward the collision point, or be retracted to a horizontal orientation to be moved aside during access periods.

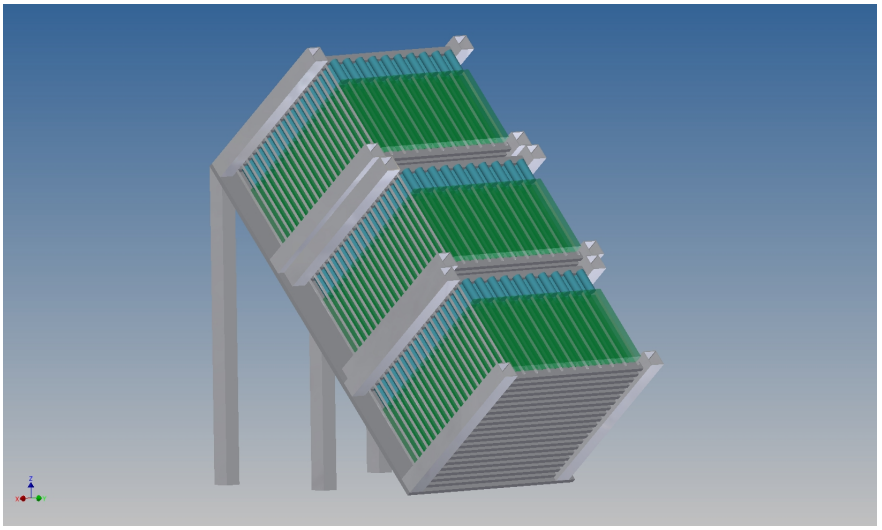

FIG. 2: A sketch of the working design for the mechanical support. The rows of scintillator bars and PMTs are mounted in trays within three separate stacks. The middle stack is offset horizontally by 1/2 unit as discussed later in this document. An adjustable platform supports all three stacks and can be tilted to point toward the collision point for data taking or be retracted to a horizontal orientation during access periods.

The sides of the assembly would be covered with a mu-metal skin for magnetic shielding and enclosed for temperature control. The size of the drainage gallery leaves little clearance at the corners of the detector, which constrains the thickness and mounting in those regions. However, the space along the sides of the detector is not constrained. We envision using

that space for the hoists necessary for positioning and retractions as well as the cable plant and cooling.

As discussed below, the dominant background is expected to arise from PMT dark current, which can be reduced by cooling the PMTs below room temperature. The limited infrastructure available in the drainage gallery hinders the use of water cooling, but we expect the estimated heat load of a few kW could be managed with forced air cooling from a locally positioned air conditioning unit.

Such a design appears to be workable without significant infrastructure demands. We are proceeding to study specific engineering options in more detail.

VI. POWER AND CALIBRATIONS

PMTs that meet the required specifications in terms of pulse rise time, dark current and counting rates, and quantum efficiency require applied high voltages (HV) between $1-2.5$ kV and have maximum current ratings of $0.2-0.5$ mA. For the detector design using 12 read-out channels per module, the HV power supply (HVPS) must provide approximately 10 mA per module. In order to minimize costs, we aim to use one HV power supply to power 10 modules (120 channels), and thus we require a HVPS rated to approximately 100 mA and 500 W. For a 100 module detector, 10 HVPS are required and the total power requirement would thus be approximately 5 kW. Several commercial HVPS systems exist that meet these requirements. For example, the Glassman model number PS/EK03R200-GK6 provides an output of ± 3 kV with a maximum of 200 mA, and features controllable constant current / constant voltage operation. Regulation and monitoring of the power supplied to the detector will be required on both the module distribution boards and the front-end distribution boards. In both cases, over-current and over-voltage protection will be necessary both for safety and in order to protect the front-end electronics from damage. The monitoring may be accomplished by a measurement circuit that digitizes and transmits the measured voltages and currents over a serial bus to the slow control system for the detector by a generic, CERN built data acquisition board called an Embedded Local Monitoring Board (ELMB) [14].

Energy calibration will be done in situ using an ^{241}Am source, which yields a 60 keV X-ray. Calibration runs performed at specified intervals will track the PMT+scintillator response as a function of time. In addition to energy calibration, an LED pulser that can

deliver a stable light pulse into each scintillator will also be deployed. The LED system will be used to monitor drift in response of the PMT+scintillator as a function of time in between ^{241}Am source calibrations as well as detect any inefficient or non-functional readout channels.

VII. TRIGGER AND READOUT

Analog pulses from the PMTs must be read out, digitized, and stored for offline analysis. Furthermore, as the pulse rate per PMT is large, a trigger will be used to record only those pulses during interesting time windows when significant activity in the detector is observed, with at least two of the three layers in a $2 \times 2 \times 3$ detector module having a pulse above the single PE threshold. Fortunately, commercial electronics are available for performing these tasks in a simple way at reasonable cost.

The triggering and readout can be performed by the CAEN V1743 digitizer, which uses the SAMLONG chip. The VME board has 16 channels, each of which is sampled at 3.2 GS/s (a sample each 312.5 ps) into a 1024 analog buffer ring (320 ns long). Analog noise is about 0.75 mV per channel, allowing good identification of and triggering on single PE signals, which can easily be above 3 mV. The only drawback of the analog buffer approach is that the board is dead while digitizing the buffer, which takes 20-125 μs, depending on how much of the buffer, i.e. what fraction of the 320 ns, is digitized. This results in a deadtime of $\lesssim 10$ %, even at 1kHz trigger rate, and the time within ~ 10 μs after a trigger is anyways of low quality due to the presence of PMT afterpulses. Seven buffers are available to hold triggered events in RAM until it is read out, either over VME or an optical link. Thus there is no dead time from readout, up to rates of ~ 1 kHz. The digitization is done with 12 bits of precision, and the dynamic range is 2.5 V, allowing good resolution of 2500 mV / 4096 = 0.61 mV.

The board is also capable of self-triggering on each channel's analog sample, via a 16-bit discriminator with a resolution of 0.04 mV, for each channel. The triggers for each pair of channels (0+1, 2+3, etc.) are then ORed (or ANDed) together within the board FPGA, so that 8 trigger bits are available per event. These are then compared, further in the board's FPGA, and can be required to form coincidences within an adjustable length time window. We envision ORing together the pairs of channels, which would be connected to channels in

the same longitudinal 2×2 detector layer. We then would ask that at least 2 such bits are fired within a 15 ns window (the smallest time window available, but roughly optimal for our pulse resolution and triggering needs). This means that at least two scintillator bars (not in the same part of a layer) are above threshold coincidently. When a coincidence is triggered in one board's channels as just discussed, it can be propagated to the other boards of the experiment via an external trigger lemo input. Data will be read out via CAEN CONET 2 over an optical fiber to a PCI card, which can sustain up to 80 Mbps, which is an order of magnitude more than the expected event data rate, as discussed in section VIII.

The CAEN V1743 digitizer board is capable of handling an external clock, and distributing it across multiple boards via an external clock sync in lemo cable. And the board has a 16-bit LVDS input that can be used as an event identifier to tag events for offline identification. Additional background rejection can be obtained by requiring scintillation signals to be coincident with the expected arrival time of particles produced in LHC collisions. This can be achieved by sending the LHC clock to the readout system and recording the phase of the clock for each triggered event. The clock signal will be sent on a fiber from the CMS Trigger and Clock Distribution System (TCDS), which will also carry CMS run start and stop signals, as well as "luminosity section" and orbit counter reset signals. These signals will allow recorded events to be correlated with the CMS luminosity database, enabling the cross-section of a signal to be measured. The TCDS fiber will be received by a dedicated clock and timing module, comprising an FPGA, an SFP carrier for reception of the optical fiber, and LVDS outputs to the readout cards. (Various FPGA evaluation boards are suitable, for example the Xilinx Spartan 6 SP605). The TCDS signals will be decoded in the FPGA, which will contain counters recording the LHC clock phase, as well as the CMS run, luminosity section and orbit number, for each trigger.

VIII. BACKGROUNDS

The dominant background is expected to come from dark-current pulses in the PMTs. Additional, sub-dominant, sources of background include activity in the scintillator from background radiation and photo-multiplier afterpulsing. In Ref. [12] we assumed a dark-current background rate for a single PMT to be 550 Hz for a single-PE threshold. To get a more accurate estimation of this rate, we constructed a test setup using a 3" Bicron-412

scintillator coupled to a 3" Hamamatsu H2431 PMT at 3 kV, readout with a CAEN V1743 digitizer, described in Section VII. Once the board receives a triggering event, the analog buffer is digitized. During this digitization window, which lasts approximately 100 μs, no additional trigger can be accepted. Using this test setup, we measured a dark-current rate of approximately 1 kHz at room temperature. By reducing the high voltage and cooling the PMT we are able to significantly reduce the background rate. The studies are still ongoing, but with these handles and by optimizing the choice of PMT we expect to be able to bring the rate to below 500 Hz.

The vast majority of pulses from background radiation, including cosmic muons, will consist of a large number (>1000) of PEs. By implementing an offline veto of these large pulses, for example those with more than 10 PEs, these backgrounds can be dramatically reduced. Whenever a pulse enters a photo-multiplier, there are smaller after-pulses that are generated. These small after-pulses occur within approximately 10 μs of the original large pulse. Since the large pulse will trigger the board, the small after-pulses will fall within the digitization deadtime of the CAEN V1743 and thus be vetoed.

The rates from all background sources will be greatly reduced by the requirement of a 3-fold coincidence between the layers of the detector. However, a remaining background source is cosmic muons which pass along the edge of all three scintillator layers. Such a glancing trajectory could result in a small pulse in each of the 3 PMTs. To account for this possibility, we will offset the middle layer, eliminating such trajectories. The remaining background from cosmic muons would be from trajectories that glance by the first and third layers and pass through the bulk of the second layer.

With the dominant background coming from the dark-current in the PMT, we assume a total background rate per PMT of ν_B =500 Hz. By reading out pairs of PMTs in the same layer, the rate per group would be $2\nu_B$. Each CAEN V1743 board would receive input from 6 such pairs, 2 in each layer, leaving 2 groups without input. With a reasonably robust time window of $(\Delta t)_{\text{online}} = 100$ ns, the double coincident trigger rate per-board will be $\binom{6}{2}(2\nu_B)^2(\Delta t)_{\text{online}} = 1.5$ Hz. The entire detector will be read out if one board triggers and there will be 50 such boards in total. Therefore, the full background trigger rate is expected to be 75 Hz. Offline we expect to be able to tighten the time window to $(\Delta t)_{\text{offline}} = 15$ ns as discussed previously. The offline background rate for a triple coincidence is given by $\nu_B^3(\Delta t)^2_{\text{offline}} = 2.8 \times 10^{-8}$ Hz. Since there are 400 such sets, the total offline background

rate is estimated to be 1.1×10^{-5} Hz. By the end of Run 3 LHC will have delivered 300 fb^{-1}. Assuming an average instantaneous luminosity of 2×10^{34} cm^{-2}s^{-1} and a 50% timing window efficiency, we calculate a trigger live-time of 1.5×10^{7} s. During HL-LHC operation, 3000 fb^{-1} will be delivered with an instantaneous luminosity 1×10^{35} cm^{-2}s^{-1}, corresponding to a a trigger live-time of 3.0×10^{7} s. We therefore estimate that we will have 165 (330) background events in 300 (3000) fb^{-1}. However, we expect to be able to use additional handles (e.g. pulse shapes of the pulses and tighter timing cuts) to further reduce the total background in these run periods to ~ 50 (100) events.

These studies provide an estimate of the background rate, which is used below to calculate the expected sensitivity. Ultimately, the background will be determined from the data, for example by measuring the rate of near signal-like events such as triple-incidence of single PE hits that do not point to the collision.

IX. SIMULATIONS AND SENSITIVITY

In order to evaluate the projected sensitivity of the experiment for various mCP electric charges and masses, we performed a full simulation of the experiment, including a GEANT4 [15] model of the detector. The simulation is performed in two stages. In the first, FEYNRULES, MADGRAPH5 and MADONIA were used to simulate the production of mCP particles via Drell-Yan, J/Ψ, Υ(1S), Υ(2S), and Υ(3S) channels at 14 TeV center-of-mass energy [16–18]. Particles produced at the interaction point were propagated using a map of the CMS magnetic field, shown in Figure 3, to the proposed experimental site described in Section II. Although small for particles with charge $Q \ll e$, the effects of multiple scattering and energy loss were included using a simplified model of the CMS detector material budget and a region of rock spanning 17 m between the CMS experimental cavern and the proposed experimental site. The number of expected mCP particles per fb^{-1} of integrated luminosity incident at the detector is shown in Figure 4 as a function of the mass of the milli-charged particle. To illustrate the dependence of the acceptance on the charge, the cross section for all charge scenarios is normalized to that for a milli-charged particle with $Q = 0.1\ e$.

In the second stage, we calculated the signal efficiency by running the kinematic distributions of the particles at the proposed experimental site through a full GEANT4 simulation of the detector, as viewed in Figure 5. This was important because although we can estimate

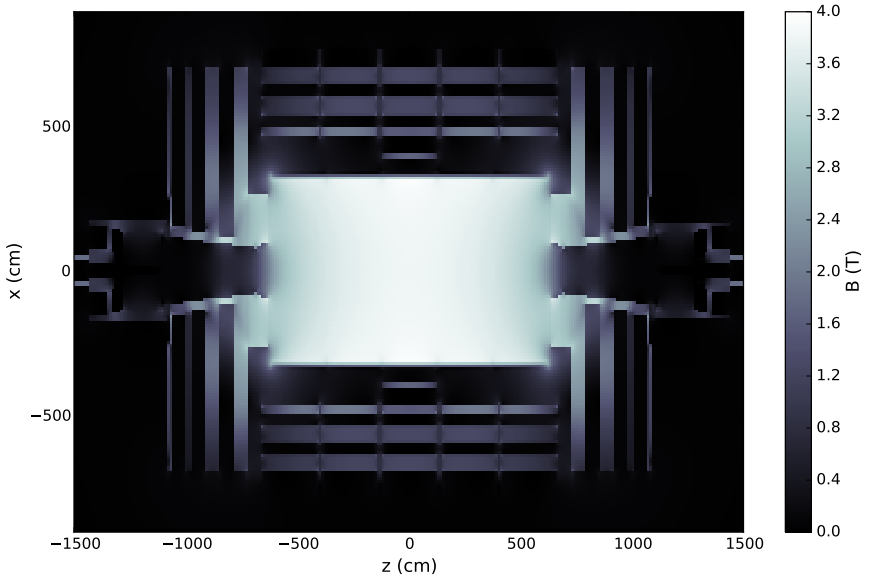

FIG. 3: Map of the CMS magnetic field in the r–z plane.

part of the efficiency of the detector by hand, the small charge regime is sensitive to details such as the reflectivity, the light attenuation length, and the shape of the scintillator. These details, as well as the quantum efficiency, light emission spectrum and the fast time constants are modeled in GEANT4 using the specifications provided by the manufacturers for Saint-Gobain BC-408 plastic scintillator and Hamamatsu R329-02 PMTs [19, 20]. We defined a new fermion of variable mass and electric charge. Its electromagnetic interactions consist of multiple scattering, Bethe-Bloch energy loss and density effects, implemented in GEANT4 using the "G4WentzelVIModel" and "G4hIonisation" packages, which are documented in the source files. Figure 6(a) shows how the efficiency of a single scintillator bar changes as a function of electric charge when varying the transverse dimensions and the reflectivity, for a 0.1 GeV mCP. Figure 6(b) shows the same plot for the full detector, requiring a 15 ns triple coincidence. In both plots, we compare the GEANT4 efficiencies to the efficiencies assumed in the Ref. [12]. The probability of seeing one or more photoelectrons in each layer of the detector was parametrized using Poisson statistics by

$$P = \left(1 - \exp\left[-N_{PE}\right]\right)^3, \tag{4}$$

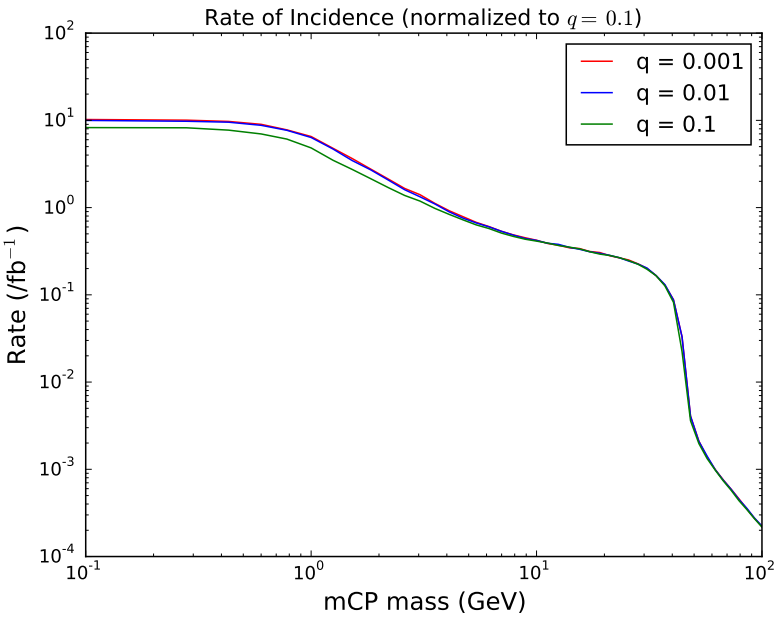

FIG. 4: Number of expected mCP particles per fb^{-1} of integrated luminosity incident at the detector as a function of the mass of the milli-charged particle. To illustrate the dependence of the acceptance on the charge, the Q^2 production dependence has been factored out by normalizing the cross section for all charge scenarios to that for a milli-charged particle with $Q = 0.1\ e$.

where $N_{PE} = \left(\frac{Q}{\xi}\right)^2$ is the average number of photoelectrons produced for a given charge. The constant of proportionality ξ was estimated by finding the electric charge that gives 1 photoelectron, given the material light yield, a 10% detection efficiency, the length of the scintillator and typical energy deposits of a minimally ionizing particle. It was found to be $\xi \approx 0.0024$. Comparing this estimate to the GEANT4 efficiencies, we find good agreement, especially for the large mass regime (not shown). The mCPs in the lower mass regime are more relativistic, and deposit less energy.

Combining the estimated background rates discussed in Section VIII with the cross-sections, acceptances and efficiencies calculated for all masses and electric charges, the sensitivity projections of the milliQan experiment for LHC and HL-LHC are shown in Figure 7.

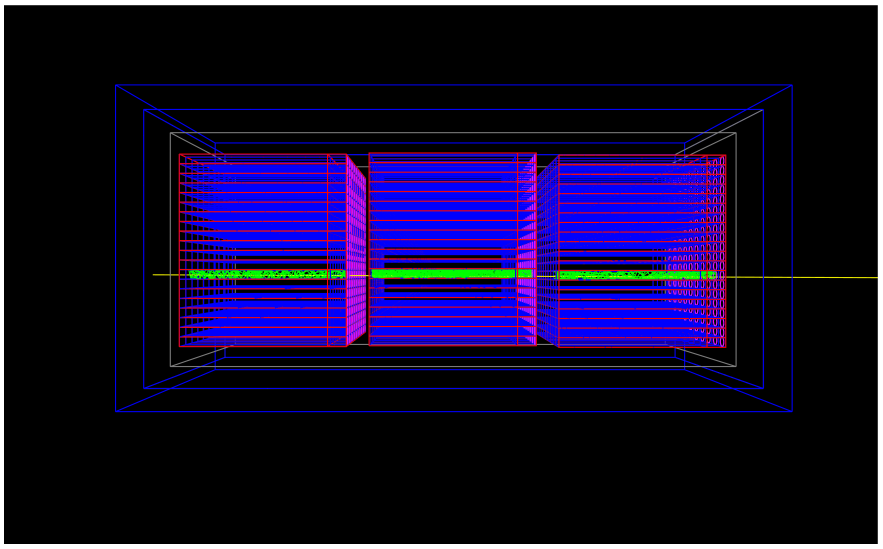

(a)

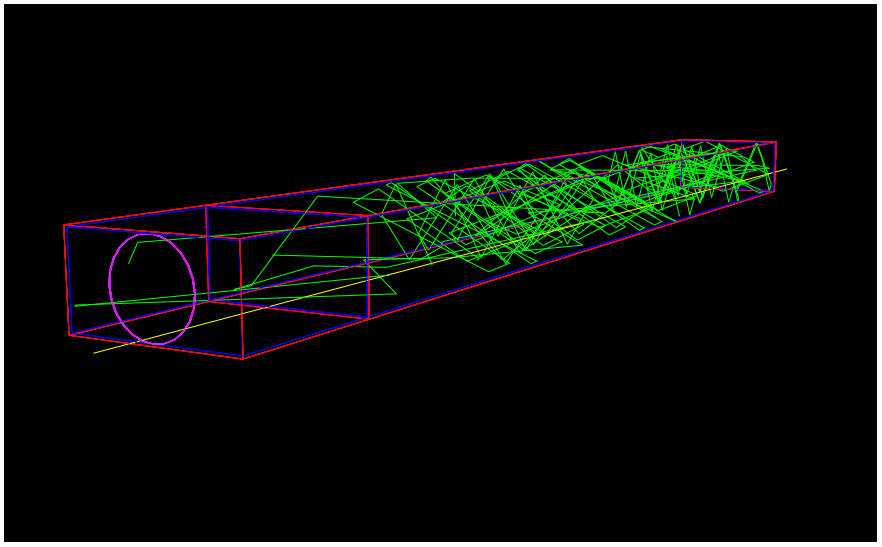

(b)

FIG. 5: Depiction of the (a) full detector and (b) a single scintillating block with coupled phototube, as implemented in the GEANT4 detector simulation. The mCP is yellow and radiated photons are green.

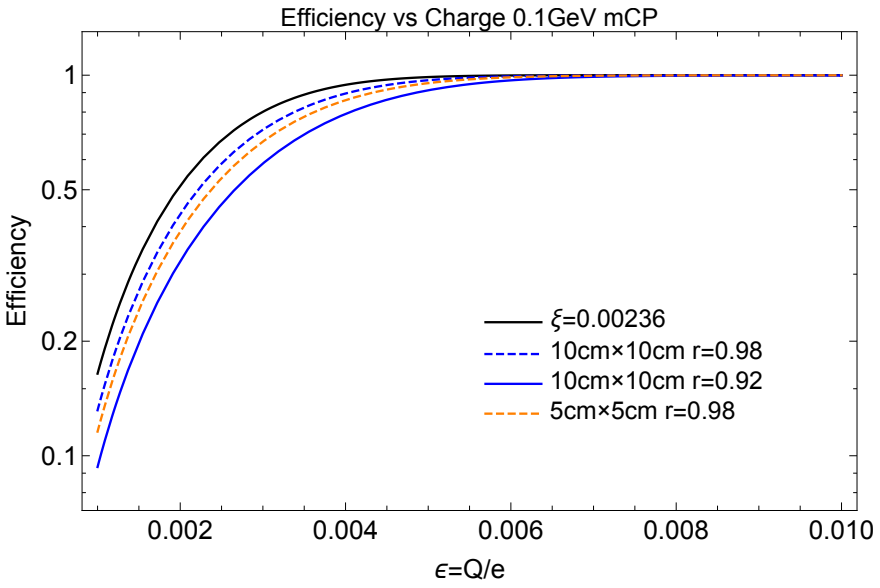

(a)

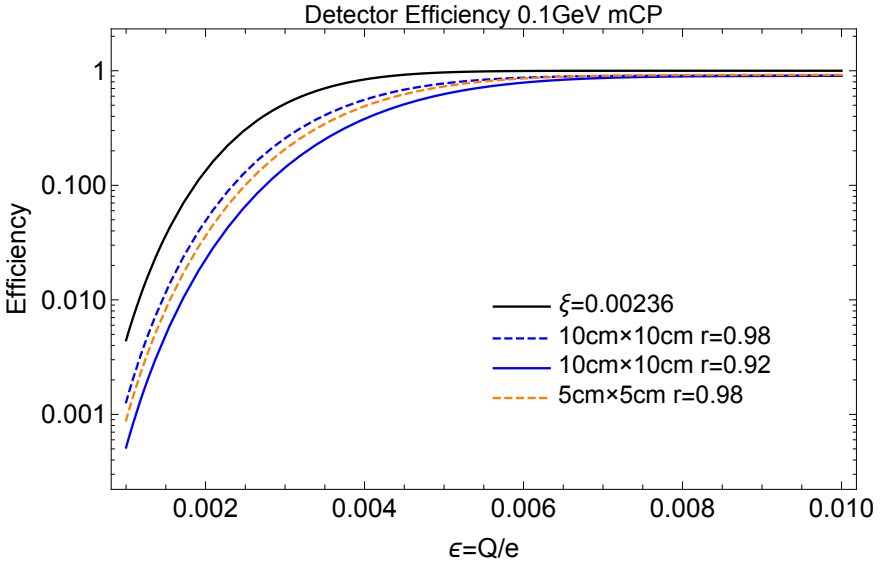

(b)

FIG. 6: Efficiencies for (a) a single scintillator block and coupled PMT and (b) the whole detector with 15ns triple-incidence, as determined from the GEANT4 detector simulation.

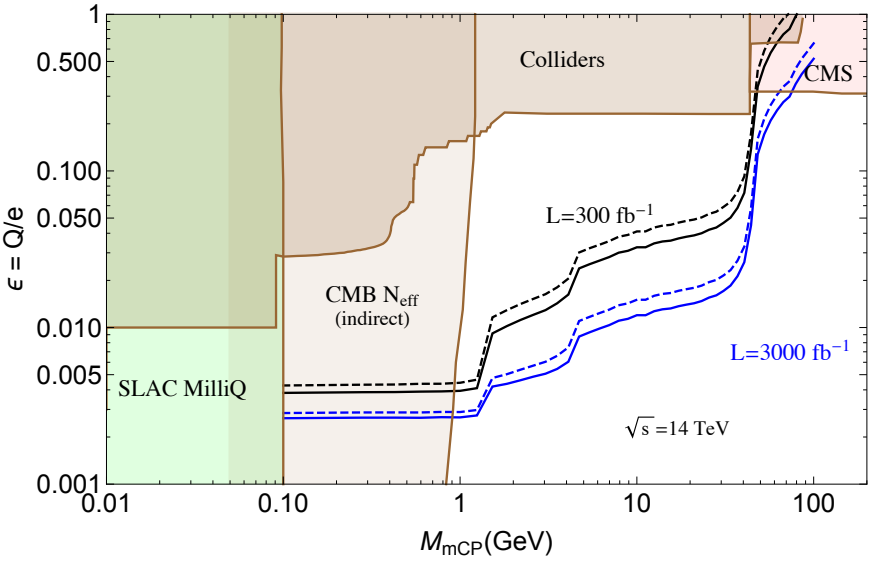

FIG. 7: Expected sensitivity for different LHC luminosity scenarios. The black line shows the expected 95% C.L. exclusion (solid) and 3σ sensitivity (dashed), assuming 300 fb^{-1} of integrated luminosity. In blue we show the corresponding expectations for 3000 fb^{-1}.

X. TIMELINE AND NEXT STEPS

We aim to have the experiment ready for physics during Run 3. To that end, we envisage the following timeline:

- Construct small fraction of detector ($\sim$ 10%) in next 2 yrs
- Install partial detector in PX56 by end of Run 2 (YETS 2017 + TS in 2018)
- Commission and take data in order to evaluate beam-on backgrounds *in situ*
- Construction + Installation of remainder of detector during LS2 (2019–2020)
- Final commissioning by spring 2021
- Operate detector for physics for duration of Run 3 and HL-LHC (mid 2021–)

The next step in the milliQan project is to seek external funding to enable at least the 10% construction. No such funding has yet been secured for this project, but one or more proposals to one or more funding agencies are being prepared for the near future.

XI. SUMMARY

In this LOI we have proposed a dedicated experiment that would detect "milli-charged" particles produced by pp collisions at LHC Point 5. The experiment would be installed during LS2 in the vestigial drainage gallery above UXC and would not interfere with CMS operations. Our calculations and simulations indicate that with 300 fb^{-1} of integrated luminosity, sensitivity to a particle with charge $\mathcal{O}(10^{-3})$ e can be achieved for masses of $\mathcal{O}(1)$ GeV, and charge $\mathcal{O}(10^{-2})$ e for masses of $\mathcal{O}(10)$ GeV. This would greatly extend the parameter space explored for particles with small charge and masses above 100 MeV. We have performed sufficient R&D to encourage us to proceed with securing funding for the project, and with this letter of intent we express the intention to do so.

Acknowledgments

We wish to thank Tiziano Camporesi, Joel Butler, and the CMS collaboration for their encouragement. We would also like to thank Vladimir Ivanchenko, Andrea Dotti and Mihaly Novak for useful discussions regarding GEANT4.

[1] B. Holdom, Phys.Lett. **B166**, 196 (1986).

[2] A. Prinz, R. Baggs, J. Ballam, S. Ecklund, C. Fertig, et al., Phys.Rev.Lett. **81**, 1175 (1998), hep-ex/9804008.

[3] S. Davidson, S. Hannestad, and G. Raffelt, JHEP **0005**, 003 (2000), hep-ph/0001179.

[4] A. Badertscher, P. Crivelli, W. Fetscher, U. Gendotti, S. Gninenko, et al., Phys.Rev. **D75**, 032004 (2007), hep-ex/0609059.

[5] S. Chatrchyan et al. (CMS Collaboration), Phys.Rev. **D87**, 092008 (2013), 1210.2311.

[6] S. Davidson, B. Campbell, and D. C. Bailey, Phys.Rev. **D43**, 2314 (1991).

[7] R. Mohapatra and I. Rothstein, Phys.Lett. **B247**, 593 (1990).

[8] S. Davidson and M. E. Peskin, Phys.Rev. **D49**, 2114 (1994), hep-ph/9310288.

[9] S. Dubovsky, D. Gorbunov, and G. Rubtsov, JETP Lett. **79**, 1 (2004), hep-ph/0311189.

[10] A. Dolgov, S. Dubovsky, G. Rubtsov, and I. Tkachev, Phys.Rev. **D88**, 117701 (2013), 1310.2376.

[11] H. Vogel and J. Redondo, JCAP **1402**, 029 (2014), 1311.2600.

[12] A. Haas, C. S. Hill, E. Izaguirre, and I. Yavin, Phys. Lett. **B746**, 117 (2015), 1410.6816.

[13] J. Beringer et al. (Particle Data Group), Phys.Rev. **D86**, 010001 (2012).

[14] H. Boterenbrood and B. I. Hallgren, p. 4 p (2003), URL `http://cds.cern.ch/record/690030`.

[15] S. Agostinelli et al. (GEANT4), Nucl. Instrum. Meth. **A506**, 250 (2003).

[16] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, JHEP **07**, 079 (2014), 1405.0301.

[17] A. Alloul, N. D. Christensen, C. Degrande, C. Duhr, and B. Fuks, Comput. Phys. Commun. **185**, 2250 (2014), 1310.1921.

[18] P. Artoisenet, F. Maltoni, and T. Stelzer, Journal of High Energy Physics **2008**, 102 (2008), URL `http://stacks.iop.org/1126-6708/2008/i=02/a=102`.

[19] Saint-Gobain, *Plastic scintillators*, `http://www.crystals.saint-gobain.com/Plastic_Scintillators.aspx`, accessed: 2016-02-12.

[20] Hamamatsu, *Photomultiplier tube r329-02*, `http://www.hamamatsu.com/us/en/product/category/3100/3001/R329-02/index.html`, accessed: 2016-02-12.

[21] Such a term is expected in grand unified theories and more generally whenever there exists massive fields that are charged under both hypercharge and the new gauge boson, even when these heavy fields are not accessible at low energies.

[22] The detector will be adequately shielded from other background sources such as activity in the scintillator and environmental radiation.

Chapter 3

Millicharged particles in neutrino experiments

3.1 Preface

In this paper, we re-interpret existing single electron measurements at MiniBooNE and LSND to derive new bounds on mCPs. We demonstrate how future experiments can improve upon these results by calculating expected single electron rates at these facilities. We focus primarily on elastic scattering with electrons, since this channel is enhanced in the low-Q^2 regime. In this limit, it is straightforward to calculate that $\sigma_{e\chi} \approx 4\pi\alpha^2\epsilon^2/Q^2_{\rm min}$, where Q is related to the recoil energy of the electron via $Q^2 = 2m_e(E_e - m_e)$. For nucleon scattering, the cross section is suppressed by $1/m_p$ rather than $1/m_e$. An experiment's recoil energy threshold, $E_e^{(\rm min)}$, then sets the scale of the detection cross section as

$$\sigma_{e\chi} = 2.6 \times 10^{-25}{\rm cm}^2 \times \epsilon^2 \times \frac{1\ \ {\rm MeV}}{E_e^{(\rm min)} - m_e}. \tag{13}$$

The data we use at LSND [87] includes electron recoils between 18 MeV$-$52 MeV. At future LAr-TPC detectors, based on their stopping power and interwire spacing of 3 mm, it may be feasible to consider electron recoils as low as 0.8 MeV $-$ 2 MeV. Consequently, sensitivity to mCPs can be greatly enhanced at future experiments, producing approximately 60 times more mCPs per POT than at

LSND. Furthermore, backgrounds can be reduced since LAr-TPC detectors can veto photons and can use accurate directional and timing information to remove other sources of backgrounds, such as cosmic particles. We consider mostly production via the π^0 and η mesons. At future experiments such as DUNE and SHiP, the POT energies are respectively 80 GeV and 400 GeV. This opens up the possibility of producing mCPs from heavy mesons such as J/Ψ and Υ, and from Drell-Yan production at the proton target. As such, these neutrino experiments may have sensitivity to mass regions that overlap with the milliQan experiment. This would provide two extremely different probes of mCPs in the mass range between 0.1 GeV and 5 GeV and would serve as a useful cross-check.

Contribution of author (reproduced from Declaration of Authorship): Ryan Plestid was responsible for upgrading the meson decay packages. Gabriel Magill was responsible for implementing/developing the pipeline to perform the sensitivity studies at the intensity frontier experiments. Yu-Dai performed the Drell-Yan calculations. Ryan and Gabriel both contributed significantly to the writing of the paper. This paper appears on arXiv as:

- Gabriel Magill, Ryan Plestid, Maxim Pospelov, and Yu-Dai Tsai, "Millicharged particles in neutrino experiments," (2018), arXiv:1806.03310 [hep-ph]

and will be submitted to PRL.

3.2 Paper

Millicharged particles in neutrino experiments

Gabriel Magill,[1,2,*] Ryan Plestid,[1,2,†] Maxim Pospelov,[1,3,4,‡] and Yu-Dai Tsai[5,§]

[1] *Perimeter Institute for Theoretical Physics, 31 Caroline St. N., Waterloo, Ontario N2L 2Y5, Canada*
[2] *Department of Physics & Astronomy, McMaster University, 1280 Main St. W., Hamilton, Ontario L8S 4M1, Canada*
[3] *Department of Physics and Astronomy, University of Victoria, Victoria, BC V8P 5C2, Canada*
[4] *Theoretical Physics Department, CERN, 1211 Geneva, Switzerland*
[5] *Laboratory for Elementary Particle Physics, Cornell University, Ithaca, NY 14850, USA*

(Dated: June 12, 2018)

We set constraints on millicharged particles (mCPs) based on electron scattering data from MiniBooNE and the Liquid Scintillator Neutrino Detector (LSND). Both experiments are found to provide new (and leading) constraints in certain mCP mass windows: $5-35$ MeV for LSND and $100-180$ MeV for MiniBooNE. Furthermore, we provide projections for the ongoing SBN program, the Deep Underground Neutrino Experiment (DUNE), and the proposed Search for Hidden Particles (SHiP) experiment. Both DUNE and SHiP are capable of probing parameter space for mCP masses ranging from 5 MeV $-$ 5 GeV that is significantly beyond the reach of existing bounds, including those from collider searches and SLAC's mQ experiment.

Introduction: The extensions of the Standard Model (SM) by light weakly charged particles, and their probes at the intensity frontier experiments have become an important direction of modern particle physics [1]. One of the simplest and most natural ways of coupling new particles to the SM is via a "kinetic mixing" or "hypercharge portal" [2, 3], which at low energy may lead to millicharged particles (mCPs), that would seemingly contradict the observed quantization of electric charge in nature [4]. In recent years, a wide class of related models were studied in connection with dark matter [5–7] (see also [8–16]), and mCPs can be viewed as a specific limit of those theories.

It is well appreciated that both proton and electron beam dump experiments provide sensitive probes of vector portal models. In particular, production and scattering of light dark matter [9] has been studied as a function of mediator mass $m_{A'}$, dark sector coupling α_D, dark matter mass m_χ, and kinetic mixing parameter ϵ_Y. Depending on the relation between these parameters, either the past electron beam dump facilities [12] or the proton fixed target experiments with a primary goal of neutrino physics [10, 13] provide the best sensitivity. However, the simplest limit of $m_{A'} \to 0$, when the parameter space simplifies to the mass and effective charge of mCPs, $\{m_\chi, \epsilon\}$, was analyzed only in the context of electron beam dump experiments [17, 18]. Clearly, fixed target neutrino experiments, such as the existing data from MiniBooNE [19] and the Liquid Scintillator Near Detector (LSND) [20], and the soon to be released data from MicroBooNE, the ongoing SBN program [21], the Deep Underground Neutrino Experiment (DUNE) [22], and the proposed Search for Hidden Particles (SHiP) [23] serve as a fertile testing ground of MeV–GeV physics due to their inherently high statistics [10, 13, 24, 25]. These experiments all serve as promising avenues to probe the mCP model.

The purpose of this *Letter* is twofold: First, we demonstrate that existing data from LSND provides leading bounds on mCPs (slightly surpassing existing constraints from SLAC's mQ experiment [17]) in the low mass regime ($m_\chi \lesssim 35$ MeV). Likewise, newly released data from MiniBooNE [19] can set more stringent bounds on mCPs in the mass range of 100 MeV $\lesssim m_\chi \lesssim$ 180 MeV. Second, we predict that by optimizing search strategies at ongoing and upcoming experiments (such as MicroBooNE,SBND, DUNE, and SHiP), fixed source neutrino experiments can serve to provide leading bounds for mCP masses over the full range of masses 5 MeV $\lesssim m_\chi \lesssim$ 5 GeV. The detection signature of mCPs in these experiments is elastic scattering with electrons, and we find that detection prospects are highly sensitive to the threshold imposed on the electron's recoil energy. Therefore, significant gains in sensitivity to mCPs may be achieved by future experiments by optimizing the detection of low energy electrons.

Our results have direct implications for models with late kinetic coupling of dark matter and baryons [30] that could lead to extra cooling of the baryon fluid and spin temperature at redshifts $z \simeq 20$, which in turn may result in a more pronounced 21 cm absorption signal. If a fraction of dark matter is in the form of mCPs, this extra cooling mechanism can be naturally realized [31, 32], and fit the unexpected strength of the signal reported by Experiment to Detect the Global Epoch of Reionization Signature (EDGES) [33]. The interpretation of the EDGES result as shedding light on dark matter-baryon interactions necessitates a careful consideration of existing laboratory constraints. In particular, our analysis reveals that sensitivities from LSND, SBND, SHiP, and DUNE can explore previously unprobed regions of parameter space that are favored by the 1%-mCP fractional dark matter hypothesis [29, 32, 34].

Production and detection: Fixed target neutrino experiments rely on the production of neutrinos from weak decays of charged pions. In generating an appro-

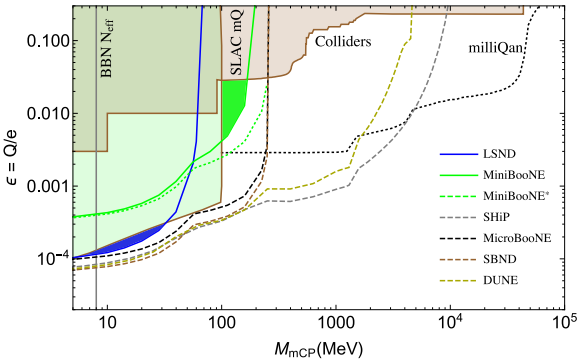

FIG. 1. Exclusion curves for fermionic mCPs (results are broadly similar for scalars). Existing data is shown as solid lines, while projections are shown as dashed curves. The kinematic reach of a given experiment is set by the heaviest meson of interest it can produce. This is π^0 for LSND, η for the Booster experiments, and Υ for DUNE. At SHiP, Drell-Yan production extends the kinematic reach to roughly 10 GeV. The sensitivity of each experiment can be understood via Eq. (4) while the relevant parameters for each experiment are summarized in Table I. The bound on $N_{\rm eff}$ [26] comes from changing the effective number of neutrinos during BBN, while the SLAC mQ and collider bounds are taken from [17] and [18, 27] respectively. The projected sensitivities at milliQan are from [27, 28]. Our exclusions apply independently of the existence of a dark photon, which would only introduce additional constraints [29].

priately large flux of $\pi^\pm$ these experiments necessarily also produce a similar number [i.e. $\mathcal{O}(10^{20})$] of π^0 [16]. For large beam energies, other neutral mesons (e.g. η, Υ, J/ψ) are also produced. Any significant branching ratios to lepton pairs necessarily implies an associated decay to pairs of mCPs, resulting in a significant flux of mCPs even for extremely small charges. In the case of η and π^0, Dalitz decays $\pi^0/\eta \to \gamma\chi\bar{\chi}$ dominate, while for J/ψ and Υ direct decays $J/\psi, \Upsilon \to \chi\bar{\chi}$ are most important. The branching ratio for a meson, $\mathcal{M}$, to mCPs is given roughly by

$$\mathrm{BR}(\mathcal{M} \to \chi\bar{\chi}) \approx \epsilon^2 \times \mathrm{BR}\big(\mathcal{M} \to X e^+ e^-\big) \times f\Big(\frac{m_\chi}{M}\Big), \quad (1)$$

where M is the mass of the parent meson, X denotes any additional particles, and $f(m_\chi/M)$ is a phase space factor that decreases slowly as a function of m_χ/M. The number of mCPs passing through the detector is a function of both the branching ratio and geometric losses which can vary substantially between experiments (see Table I).

We now turn to the detection of mCPs at neutrino beam dump experiments, where the predominant signature is elastic scattering with electrons. The dominance of electron scattering as a detection signal is related to the low-Q^2 sensitivity of the scattering cross section. Explicitly, in the limit of small electron mass, we have

$$\frac{\mathrm{d}\sigma_{e\chi}}{\mathrm{d}Q^2} = 2\pi\alpha^2\epsilon^2 \times \frac{2(s-m_\chi^2)^2 - 2sQ^2 + Q^4}{(s-m_\chi^2)^2 Q^4}. \quad (2)$$

Upon integrating over momentum transfers, we see that the total cross section will be dominated by the small-Q^2 contribution to the integral. In this limit, we have $\mathrm{d}\sigma_{e\chi}/\mathrm{d}Q^2 \approx 4\pi\alpha^2\epsilon^2/Q^4$, and so we can see immediately that $\sigma_{e\chi} \approx 4\pi\alpha^2\epsilon^2/Q^2_{\rm min}$. We may relate $Q_{\rm min}$ in the lab frame to the recoil energy of the electron via $Q^2 = 2m_e(E_e - m_e)$ [35]. An experiment's recoil energy threshold, $E_e^{(\rm min)}$, then sets the scale of the detection cross section as

$$\sigma_{e\chi} = 2.6 \times 10^{-25}\mathrm{cm}^2 \times \epsilon^2 \times \frac{1\ \mathrm{MeV}}{E_e^{(\rm min)} - m_e}. \quad (3)$$

Consequently, sensitivity to mCPs can be greatly enhanced by accurately measuring low electron energy recoils (an important feature for search strategies at future experiments).

Results: We now discuss the details of the modelling and analysis used to create Fig. 1. The various curves are obtained by performing a sensitivity analysis [36]: given a number of predicted background events b and data n, the number of signal events $s_{\rm up}$ consistent with the observation and backgrounds at $(1-\alpha)$ credibility level is found by solving the equation $\alpha = \Gamma(1+n, b+s_{\rm up})/\Gamma(1+n, b)$ where $\Gamma(x,y)$ is the upper incomplete gamma function [37]. Throughout this paper, we choose a credibility interval of $1-\alpha = 95\%$ and calculate the corresponding bounds implied by $s_{\rm up}$ on our mCP model according to the formula

$$s_{\rm up} = \sum_{\rm Energies} \epsilon^4 \times N_\chi(E_i) \times \frac{N_e}{\rm Area} \times \sigma_{e\chi}(E_i;\, m_\chi) \times \mathcal{E}. \quad (4)$$

Here, ϵ is the mCP electric charge (in units of e), $N_\chi(E_i)$ represents the number of mCPs with energy E_i arriving at the detector, $\sigma_{e\chi}(E_i)$ is the detection cross section consistent with the angular and recoil cuts in the experiment, N_e is the total number of electrons inside the active volume of the detector, $\mathcal{E}$ is an overall electron detection efficiency. Finally, "Area" in (4) stands for the active volume divided by the average length $\langle l \rangle$ traversed by particles inside the detector. The total exposure is contained in $N_\chi(E_i)$. For most of the mCP parameter space under consideration, electromagnetic decays of mesons provide the dominant flux contribution, whereas Drell-Yan production (DYP) dominates for the large mCP masses that are only accessible at DUNE and SHiP.

To estimate how many mCPs of energy E_i arrive at the detector, we model the angular and energy distributions of the mesons using one of several empirical formulas to be discussed below. Given a meson produced at a certain angle and energy, we numerically sample its branching

	N [$\times 10^{20}$]		$A_{\rm geo}(m_\chi)[\times 10^{-3}]$		Cuts [MeV]		
Exp.	π^0	η	1 MeV	100 MeV	$E_e^{\rm min}$	$E_e^{\rm max}$	Bkg
LSND	130	—	20	—	18	52	300
mBooNE	17	0.56	1.2	0.68	130	530	2K
mBooNE*	1.3	0.04	1.2	0.68	18	—	0^*
μBooNE	9.2	0.31	0.09	0.05	0.8	40	16
SBND	4.6	0.15	4.6	2.6	0.8	40	240
DUNE	830	16	3.3	5.1	2	40	$19K$
SHiP	4.7	0.11	130	220	20	50	25

TABLE I. Summary of the lifetime meson rates (N), mCP detector acceptances ($A_{\rm geo}$), electron recoil energy cuts, and backgrounds at each of the experiments considered in this paper. In all experiments a cut of $\cos\theta > 0$ is imposed in our analysis (*except for at MiniBooNE's dark matter run where a cut of $\cos\theta > 0.99$ effectively reduces backgrounds to zero [38]). For the SHiP and DUNE experiments, we also include J/ψ and Υ mesons as well as Drell-Yan production which are discussed in the text. We use an efficiency of $\mathcal{E} = 0.2$ for Cherenkov detectors, $\mathcal{E} = 0.5$ for nuclear emulsion detectors, and $\mathcal{E} = 0.8$ for liquid argon time projection chambers. The data at LSND and MiniBooNE is taken from [39] and [19] respectively. Projections at MiniBooNE* [40], MicroBooNE [41], SBND [21], DUNE [22] and SHiP [42] are based on expected detector performance.

ratio to mCPs over all possible angles and energy in the lab frame, and determine the fraction of its branching ratio to mCPs in which one of such particles has energy E_i and is pointed towards the detector. Repeating this procedure over all production energies and angles of the meson yields the meson contribution to $N_\chi(E_i)$. For DYP of mCPs from a quark and anti-quark pair, we integrate over the full production phase-space using MSTW parton distribution functions [43], and using Heaviside functions, we select the proportion of events containing an mCP pointed towards the detector, with energy E_i.

Having given a general overview of how our sensitivities are obtained, we now focus the discussion on the details of each experiment. In Table I, we show for each experiment: the lifetime rates for π^0 and η mesons, the geometric acceptance $A_{\rm geo}(m_\chi)$ [44], the cuts that we have imposed, and the expected number of background events. Using Eq. (4) this is sufficient information to approximately reproduce our results.

At LSND, the π^0 spectrum is modelled using a Burman-Smith distribution [45, 46] assuming 2 years of operation on a water target and 3 years of operation on a tungsten target. Our LSND analysis is based on [39], which featured 1.7×10^{23} protons on target (POT), a beam energy of 0.798 GeV, and a single electron background of approximately 300 events with energies ranging between 18 MeV and 52 MeV. We estimate the N_e/Area in Eq. (4) to be $2.5 \times 10^{26}\ e^-/\text{cm}^2$.

The resultant meson spectrum from Fermilab's Booster Neutrino Beam (BNB) is relevant for MiniBooNE, MicroBooNE, and SBND. The BNB delivers 8.9 GeV protons on target and so can produce substantial numbers of both π^0 and η mesons. The former's angular and energy spectra are modelled by the Sanford-Wang distribution [16, 47], and η mesons by the Feynman Scaling hypothesis [47]. These distributions are common across all three of the aforementioned experiments. We have compared our geometric acceptances with those generated using [16] and reasonable (to within an $\mathcal{O}(1)$ factor) agreement.

At MiniBooNE we perform two distinct analyses: First we consider the recently updated neutrino oscillation search [19]. We combine data from both neutrino and anti-neutrino runs and consider a sample of 2.41×10^{21} POT for which we take the single electron background to be 2.0×10^3 events and the measured rate to be 2.4×10^3. Next, motivated by a dedicated dark matter search with 1.86×10^{20} protons on target [48], we consider an anticipated parallel analysis [40] involving electron-recoil data. Backgrounds were suppressed by operating the beamline in an "off-target" mode, (i.e. not collimating charged pions), and these can be further suppressed (to zero) by imposing a cut of $\cos\theta > 0.99$ on the electron's recoil angle [38]. In both cases we estimate an electron number density of $3.2 \times 10^{26}\ e^-/\text{cm}^2$. The sensitivity curve quoted in Fig. 1 assumes that the upcoming analysis reports no signal consistent with mCPs.

At MicroBooNE, the meson rates assume 1.32×10^{21} POT and we estimate that the detector has an electron density of $3.9 \times 10^{26}\ e^-/\text{cm}^2$. The chosen recoil cuts are based on the lowest reaches achievable given the wire spacing in MicroBooNE's liquid argon detector [41]. The wire spacing is 3 mm and the ionization stopping power is approximately 2.5 MeV/cm, so electrons with energy larger than 0.8 MeV produce tracks long enough to be reconstructed. Based on this and the requirement for ionization signals that don't shower, we limit ourselves to recoil cuts between 0.8 MeV and 40 MeV. The treatment of SBND is broadly similar to MicroBooNE, but we assume 6.6×10^{20} POT, which corresponds to half the run time of MicroBooNE.

At SHiP our results assume 2×10^{20} POT and a near detector 50 m from the beam stop with an electron density of $2.7 \times 10^{26}\ e^-/\text{cm}^2$. The large beam energies of 400 GeV allow us to include J/ψ and Υ, in addition to π^0 and η. We do not include mesons such as ρ, ω and ϕ, because they do not serve to significantly alter the sensitivity offered by J/ψ (although their inclusion would only serve to increase sensitivity at SHiP for $m_\chi \lesssim 400$ MeV). At the energies of SHiP, production of π^0 and η can be described by the BMPT distribution [16, 49]. These distributions are slightly different depending on the mass of the meson with the η having a spectrum that is more forward pointed. We have compared our geometric acceptances to those obtained using [16] and found reason-

able agreement, with our acceptances being smaller by a factor of four. For production of J/ψ, we assume that their energy production spectra are described by the distributions in [50]. These distributions rely on production being highly peaked in the forward direction and parameterized as $\mathrm{d}\sigma/\mathrm{d}x_F \propto (1-|x_F|)^5$, where $x_F = 2p_\parallel/\sqrt{s}$ is the meson's longitudinal component in the COM frame of the collision. We account for geometric losses by using an empirical formulae for the p_T distribution provided in [51]. We assume that the production spectrum of Υ mesons are similarly given, and normalize their total cross section to the data in [52]. Using this, we have reproduced the Pb rates in Table 3 of [53] for J/Ψ, and for Υ we reproduced the Pt rates in Table 1 of [54]. As for our results in Fig. 1, we estimate $N_{J/\psi} = 2.1 \times 10^{15}$ with an acceptance of $A_{\rm geo}(100\text{ MeV}) = 8 \times 10^{-2}$, and $N_\Upsilon = 1.2 \times 10^{11}$ with $A_{\rm geo}(100\text{ MeV}) = 7.2 \times 10^{-2}$. For large mCP masses, DYP becomes the main production mechanism. We calibrate our DYP calculations by reproducing the dimuon invariant mass spectrum in Fig. 11 of [55] from the FNAL-772 experiment [56].

At DUNE, our treatment of meson production is very similar to the treatment at SHiP. We model pseudoscalar meson production using the BMPT distribution, as before, but use a beam energy of 80 GeV [22] and account for differences in the target material. We also include J/ψ and Υ mesons and treat them as described above. Our detector treatment and electron recoil cuts are motivated by the capability of MicroBooNE's liquid argon time projection chamber (LAr-TPC) detector, and in particular its ability to measure low energy electron recoils. We assume 3×10^{22} POT and a 30 tonne liquid argon detector which corresponds to $5.4 \times 10^{25}\ e^-/\mathrm{cm}^2$. We estimate $N_{J/\psi} = 3 \times 10^{16}$ with an acceptance of $A_{\rm geo}(100\text{ MeV}) = 2.4 \times 10^{-3}$ and $N_\Upsilon = 5.1 \times 10^9$ with $A_{\rm geo}(100\text{ MeV}) = 3.7 \times 10^{-3}$. Lastly, it is important to point out that our results do not include multiple scattering effects through dirt. Low velocity mCPs with a moderate charge (i.e. $\epsilon \gtrsim 0.03$) might get impeded by their long transit through dirt. This is relevant for DYP at DUNE and could weaken our sensitivity for $m_\chi \gtrsim 2$ GeV. Larger ϵ may also lead to a double scattering of mCPs inside the detectors, which could be used as an additional tool of discriminating their signature against the neutrino background.

We now discuss our modelling of the single electron backgrounds appearing in Table I. We consider two classes of backgrounds: those coming from each experiments flux of neutrinos [i.e. $\nu e \to \nu e$ and $\nu n \to ep$], and those coming from external sources such as cosmics, mis-identified particles, or dirt related events.

We treat neutrino induced backgrounds in detail for each experiment by summing over the neutrino fluxes provided by each collaboration and accounting for the detection efficiencies $\mathcal{E}$. Furthermore, a large background reduction is obtained by imposing the electron recoil cuts $E_e^{(\max)}$ shown in Table I. These do not significantly affect the signal (which is dominated by low electron recoils), but significantly reduce charged and neutral current backgrounds [57, 58].

We model the external sources of backgrounds by multiplying the neutrino induced backgrounds by an overall multiplicative factor. LAr-TPC detectors can use timing and directionality information as vetoes to reduce additional sources of backgrounds; this is not possible in a nuclear emulsion chamber. Therefore, we multiply our neutrino induced backgrounds by a factor of 10 for LAr-TPC detectors (MiniBooNE, SBND, and DUNE) and a factor of 25 for nuclear emulsion detectors (SHiP); this increase in the backgrounds decreases our sensitivity to ϵ by $20-30\%$. Although our naive procedure likely overestimates the backgrounds, we emphasize that our results in Fig. 1 can be easily revised for different background assumptions according to [37].

Outlook: We have shown that millicharged particles can be effectively probed at fixed target neutrino experiments due to large number of mesons produced with electromagnetic decay pathways. This includes using existing data from both LSND and MiniBooNE that are able to provide the leading sensitivity to mCPs for certain sub-GeV masses. Beyond serving as a probe of fundamental physics questions such as charge quantization, this newfound sensitivity has implications for models of physics beyond the Standard Model. In particular it further restricts the parameter space of cosmological models where a fraction of mCP dark matter results in extra cooling of baryons that modifies 21 cm physics at high redshifts.

Equally important are our projected sensitivities at MicroBooNE, SBND, DUNE and SHiP. The successful deployment of these experiments as probes of mCPs will rely heavily on their respective collaboration's search strategy. In particular by working to increase the sensitivity to low energy electron recoils the predicted signal rate can be enhanced, with a scaling proportional to $1/(E_e - m_e)$. MicroBooNE in particular has shown preliminary work that suggests good sensitivity to electron recoils with kinetic energies as low as 300 keV is possible [41]. If this can be achieved, it is conceivable that the combined sensitivity of LSND, SBND, MicroBooNE, and SHiP could provide the leading sensitivity to mCPs in the full range of 5 MeV $\lesssim m_\chi \lesssim$ 5 GeV.

Finally, we close by noting that besides the discussed current and future neutrino experiments, further progress may come from new experimental concepts. Significant progress may come from coupling large underground neutrino detectors with purposely installed new accelerators [13, 59]. Millicharged particles may also be searched by experiments in *disappearance* channels [60–62], where $e^+e^- \to \gamma + \chi + \bar{\chi}$ and $Z + e^- \to Z + e^- + \chi + \bar{\chi}$ production leads to anomalous missing momentum/energy from the χ-pair that pass through a detector without

depositing energy. Because of the advantageous scaling with ϵ (second, rather than the fourth power), there are clear prospects on improving bounds on mCPs above the 100 MeV energy range.

Acknowledgements: We thank Patrick deNiverville, Ornella Palamara, and Zarko Pavlovic for helping us to understand detector capabilities at Mini- and MicroBooNE. We would also like to thank Dan Hooper, Matthew Klimek, Gordan Krnjaic, and Tracy Slatyer for useful discussions. RP and GM are supported by the National Science and Engineering Research Council of Canada (NSERC). YT is supported in part by the U.S. National Science Foundation through grant PHY-1316222. This research was supported in part by Perimeter Institute for Theoretical Physics. Research at Perimeter Institute is supported by the Government of Canada through the Department of Innovation, Science and Economic Development and by the Province of Ontario through the Ministry of Research and Innovation.

* gmagill@perimeterinstitute.ca
† plestird@mcmaster.ca
‡ mpospelov@perimeterinstitute.ca
§ yt444@cornell.edu

[1] M. Battaglieri *et al.*, (2017), arXiv:1707.04591 [hep-ph].
[2] B. Holdom, Phys. Lett. **166B**, 196 (1986).
[3] E. Izaguirre and I. Yavin, Phys. Rev. **D92**, 035014 (2015), arXiv:1506.04760 [hep-ph].
[4] M. I. Dobroliubov and A. Yu. Ignatiev, Phys. Rev. Lett. **65**, 679 (1990).
[5] D. E. Brahm and L. J. Hall, Phys. Rev. **D41**, 1067 (1990).
[6] C. Boehm and P. Fayet, Nucl. Phys. **B683**, 219 (2004), arXiv:hep-ph/0305261 [hep-ph].
[7] M. Pospelov, A. Ritz, and M. B. Voloshin, Phys. Lett. **B662**, 53 (2008), arXiv:0711.4866 [hep-ph].
[8] J. D. Bjorken, R. Essig, P. Schuster, and N. Toro, Phys. Rev. **D80**, 075018 (2009), arXiv:0906.0580 [hep-ph].
[9] B. Batell, M. Pospelov, and A. Ritz, Phys. Rev. **D80**, 095024 (2009), arXiv:0906.5614 [hep-ph].
[10] P. deNiverville, M. Pospelov, and A. Ritz, Phys. Rev. **D84**, 075020 (2011), arXiv:1107.4580 [hep-ph].
[11] E. Izaguirre, G. Krnjaic, P. Schuster, and N. Toro, Phys. Rev. **D88**, 114015 (2013), arXiv:1307.6554 [hep-ph].
[12] B. Batell, R. Essig, and Z. Surujon, Phys. Rev. Lett. **113**, 171802 (2014), arXiv:1406.2698 [hep-ph].
[13] Y. Kahn, G. Krnjaic, J. Thaler, and M. Toups, Phys. Rev. **D91**, 055006 (2015), arXiv:1411.1055 [hep-ph].
[14] B. A. Dobrescu and C. Frugiuele, JHEP **02**, 019 (2015), arXiv:1410.1566 [hep-ph].
[15] P. Coloma, B. A. Dobrescu, C. Frugiuele, and R. Harnik, JHEP **04**, 047 (2016), arXiv:1512.03852 [hep-ph].
[16] P. deNiverville, C.-Y. Chen, M. Pospelov, and A. Ritz, Phys. Rev. **D95**, 035006 (2017), arXiv:1609.01770 [hep-ph].
[17] A. A. Prinz *et al.*, Phys. Rev. Lett. **81**, 1175 (1998), arXiv:hep-ex/9804008 [hep-ex].
[18] S. Davidson, S. Hannestad, and G. Raffelt, JHEP **05**, 003 (2000), arXiv:hep-ph/0001179 [hep-ph].
[19] A. A. Aguilar-Arevalo *et al.* (MiniBooNE), (2018), arXiv:1805.12028 [hep-ex].
[20] C. Athanassopoulos *et al.* (LSND), Nucl. Instrum. Meth. **A388**, 149 (1997), arXiv:nucl-ex/9605002 [nucl-ex].
[21] M. Antonello *et al.* (LAr1-ND, ICARUS-WA104, MicroBooNE), (2015), arXiv:1503.01520 [physics.ins-det].
[22] R. Acciarri and *et. al.*, **2** (2015), arXiv:1512.06148.
[23] M. Anelli *et al.* (SHiP), (2015), arXiv:1504.04956 [physics.ins-det].
[24] M. Pospelov and Y.-D. Tsai, (2017), arXiv:1706.00424 [hep-ph].
[25] G. Magill, R. Plestid, M. Pospelov, and Y.-D. Tsai, (2018), arXiv:1803.03262 [hep-ph].
[26] C. Boehm, M. J. Dolan, and C. McCabe, JCAP **1308**, 041 (2013), arXiv:1303.6270 [hep-ph].
[27] A. Haas, C. S. Hill, E. Izaguirre, and I. Yavin, Phys. Lett. **B746**, 117 (2015), arXiv:1410.6816 [hep-ph].
[28] A. Ball *et al.*, (2016), arXiv:1607.04669 [physics.ins-det].
[29] A. Berlin, D. Hooper, G. Krnjaic, and S. D. McDermott, (2018), arXiv:1803.02804 [hep-ph].
[30] H. Tashiro, K. Kadota, and J. Silk, Phys. Rev. **D90**, 083522 (2014), arXiv:1408.2571 [astro-ph.CO].
[31] R. Barkana, Nature **555**, 71 (2018), arXiv:1803.06698 [astro-ph.CO].
[32] J. B. Muñoz and A. Loeb, (2018), arXiv:1802.10094 [astro-ph.CO].
[33] J. D. Bowman, A. E. E. Rogers, R. A. Monsalve, T. J. Mozdzen, and N. Mahesh, Nature **555**, 67 (2018).
[34] R. Barkana, N. J. Outmezguine, D. Redigolo, and T. Volansky, (2018), arXiv:1803.03091 [hep-ph].
[35] Note that for nucleon scattering, the cross section $\sigma \propto 1/Q^2$ is suppressed by $1/m_p$ rather than $1/m_e$.
[36] C. Amsler *et al.* (Particle Data Group), Physics Letters **B667**, 1 (2008 and 2009 partial update for the 2010 edition).
[37] Given Eq. (4) and $\alpha = \Gamma(1+n, b+s_{\rm up})/\Gamma(1+n, b)$, any of our limits on ϵ can be rescaled for different choices of backgrounds. This is important for our projected sensitivities since detailed modelling of backgrounds will likely differ from what is assumed here.
[38] R. Dharmapalan *et al.* (MiniBooNE), (2012), arXiv:1211.2258 [hep-ex].
[39] L. B. Auerbach *et al.* (LSND), Phys. Rev. **D63**, 112001 (2001), arXiv:hep-ex/0101039 [hep-ex].
[40] We anticipate the release of electron recoil data from MiniBooNE's dark matter run. Our modelling of backgrouds is motivated by their nucleon scattering data [48].
[41] R. Acciarri *et al.* (MicroBooNE), JINST **12**, P09014 (2017), arXiv:1704.02927 [physics.ins-det].
[42] S. Alekhin *et al.*, Rept. Prog. Phys. **79**, 124201 (2016), arXiv:1504.04855 [hep-ph].
[43] A. D. Martin, W. J. Stirling, R. S. Thorne, and G. Watt, Eur. Phys. J. **C63**, 189 (2009), arXiv:0901.0002 [hep-ph].
[44] Defined as the ratio between the number of mCPs that reach the detector and the total number produced.
[45] R. Burman and E. Smith, (1989), 10.2172/6167579.
[46] R. Burman, M. Potter, and E. Smith, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **291**, 621 (1990).
[47] A. A. Aguilar-Arevalo *et al.* (MiniBooNE Collaboration), Phys. Rev. D **79**, 072002 (2009).
[48] A. A. Aguilar-Arevalo *et al.* (MiniBooNE), Phys. Rev. Lett. **118**, 221803 (2017), arXiv:1702.02688 [hep-ex].

[49] M. Bonesini, A. Marchionni, F. Pietropaolo, and T. Tabarelli de Fatis, Eur. Phys. J. **C20**, 13 (2001), arXiv:hep-ph/0101163 [hep-ph].
[50] C. Gale, S. Jeon, and J. I. Kapusta, Phys. Lett. **B459**, 455 (1999), arXiv:nucl-th/9812056 [nucl-th].
[51] A. Gribushin *et al.* (E672, E706), Phys. Rev. **D62**, 012001 (2000), arXiv:hep-ex/9910005 [hep-ex].
[52] I. Abt *et al.* (HERA-B), Phys. Lett. **B638**, 13 (2006), arXiv:hep-ex/0603015 [hep-ex].
[53] B. Alessandro *et al.* (NA50), Eur. Phys. J. **C48**, 329 (2006), arXiv:nucl-ex/0612012 [nucl-ex].
[54] S. W. Herb *et al.*, Phys. Rev. Lett. **39**, 252 (1977).
[55] W. J. Stirling and M. R. Whalley, J. Phys. **G19**, D1 (1993).
[56] D. M. Alde *et al.*, Phys. Rev. Lett. **64**, 2479 (1990).
[57] A. Strumia and F. Vissani, Phys. Lett. **B564**, 42 (2003), arXiv:astro-ph/0302055 [astro-ph].
[58] G. Radel and R. Beyer, Modern Physics Letters A **08**, 1067 (1993).
[59] E. Izaguirre, G. Krnjaic, and M. Pospelov, Phys. Rev. **D92**, 095014 (2015), arXiv:1507.02681 [hep-ph].
[60] J. P. Lees *et al.* (BaBar), Phys. Rev. Lett. **119**, 131804 (2017), arXiv:1702.03327 [hep-ex].
[61] D. Banerjee *et al.* (NA64), Phys. Rev. Lett. **118**, 011802 (2017), arXiv:1610.02988 [hep-ex].
[62] E. Izaguirre, G. Krnjaic, P. Schuster, and N. Toro, Phys. Rev. **D91**, 094026 (2015), arXiv:1411.1404 [hep-ph].

Chapter 4

Dipole portal to heavy neutral leptons

4.1 Preface

In this paper, we consider a phenomenological search for new physics based on single photon signatures. The model we consider, HNLs, consists of a new heavy neutral lepton coupling to neutrinos and photons via a dimension 5 dipole operator. We also consider variations of this model that make it symmetric under the symmetries of the SM. One of the motivations for this model, besides its simplicity, is that it is put forth as a possible joint explanation of the MiniBooNE and LSND excesses. This region of interest is calculated by considering production of HNLs via mass mixing, and decays to single photon final states via an enhanced dipole operator. We show that existing data can actually constrain such explanations, based on both production and decay via the dipole operator.

Furthermore, it is the case that this model has received relatively little attention outside the couplings and mass ranges relevant for the excesses mentioned above. As such, we perform a comprehensive survey across a very large range of masses and couplings. For heavier masses, we recast searches at the LHC and LEP colliders, focusing on single photon signatures without and with an accompanying lepton (this distinction is relevant for the electroweak complete versions

of this model). At very low couplings, we derive bounds coming from SN1987A cooling. In a certain region of parameter space, the HNLs can be produced in sufficiently large quantities inside the supernova and is weakly interacting enough to escape the supernova. The combination of these two effects would lead to enhanced modes of cooling, and would alter the observed neutrino signal that was detected on earth.

This paper appears on arXiv as

- Gabriel Magill, Ryan Plestid, Maxim Pospelov, and Yu-Dai Tsai, "Dipole portal to heavy neutral leptons," (2018), arXiv:1803.03262 [hep-ph]

and has been submitted for publication in Phys. Rev. D.

Contribution of author (reproduced from Declaration of Authorship): Ryan Plestid significantly developed the meson decay packages and optical depth calculations, and wrote the sections pertaining to supernovae as well as other sections of the paper. Maxim Pospelov entirely wrote the BBN section as well as parts of the introduction and conclusion. Gabriel Magill worked primarily on all of the other calculations that appear in the paper. These include doing all the sensitivity studies at LEP, LHC, MiniBooNE, MicroBooNE, SBND, SHiP, LSND, doing the coherent, diffractive and Drell-Yan calculations, coding some of the meson production phase space distributions (in association with Ryan) and writing the sections pertaining to these topics. Gabriel also contributed significantly to the emissivity and optical depth calculations of SN1987A.

4.2 Paper

Dipole portal to heavy neutral leptons

Gabriel Magill,[1, 2, *] Ryan Plestid,[1, 2, †] Maxim Pospelov,[1, 3, 4, ‡] and Yu-Dai Tsai[5, §]

[1] *Perimeter Institute for Theoretical Physics, 31 Caroline St. N., Waterloo, Ontario N2L 2Y5, Canada*
[2] *Department of Physics & Astronomy, McMaster University, 1280 Main St. W., Hamilton, Ontario L8S 4M1, Canada*
[3] *Department of Physics and Astronomy, University of Victoria, Victoria, BC V8P 5C2, Canada*
[4] *Theoretical Physics Department, CERN, 1211 Geneva, Switzerland*
[5] *Fermilab, Fermi National Accelerator Laboratory, Batavia, IL 60510, USA*

(Dated: September 12, 2018)

We consider generic neutrino dipole portals between left-handed neutrinos, photons, and right-handed heavy neutral leptons (HNL) with Dirac masses. The dominance of this portal significantly alters the conventional phenomenology of HNLs. We derive a comprehensive set of constraints on the dipole portal to HNLs by utilizing data from LEP, LHC, MiniBooNE, LSND as well as observations of Supernova 1987A and consistency of the standard Big Bang Nucleosynthesis. We calculate projected sensitivities from the proposed high-intensity SHiP beam dump experiment, and the ongoing experiments at the Short-Baseline Neutrino facility at Fermilab. Dipole mediated Primakoff neutrino upscattering and Dalitz-like meson decays are found to be the main production mechanisms in most of the parametric regime under consideration. Proposed explanations of LSND and MiniBooNE anomalies based on HNLs with dipole-induced decays are found to be severely constrained, or to be tested in the future experiments.

I. INTRODUCTION

The Standard Model of particles and fields (SM) shows remarkable resilience under the scrutiny of numerous particle physics experiments. In particular, the LHC experiments have put significant constraints on new hypothetical colored states, pushing their masses to a TeV scale and beyond. At the same time, owing to its smaller production cross sections, the electroweak extensions of the SM are far less constrained, and a plethora of new models may be hiding at energies of a few hundred GeV and beyond. If such sectors are considered to be heavy, their impact on the SM physics can be encoded in the higher-dimensional extensions of the SM. Moreover, the electroweak singlet components of such sectors can be light, and still coupled to the SM states. In the last few years, significant attention has been paid to the models containing new singlet fermionic states N (often referred to as heavy neutral leptons) that can couple to the SM leptons L and Higgs field H via the so-called neutrino portal coupling, NLH (see *e.g.* [1, 2]). Owing to the neutrality of N, its mass m_N is a free parameter with a wide range of possibilities from the sub-eV scale and up, all the way to the Planck scale. This range is somewhat narrower if N is indeed taking part in generating masses for the light active neutrino species. A great deal of experimental activity is devoted to searches of N particles, that may show up in cosmological data, in neutrino oscillation experiments, in meson decays, beam dump experiments and at high energy colliders. (For a recent overview of neutrino portal see *e.g.* [3].)

Given large interests in searches of heavy neutral leptons, in this work we will analyze a less conventional case of N particles coupled to the SM via the so-called dipole portal encoded in the following effective Lagrangian,

$$\mathcal{L} \supset \bar{N}(i\partial\!\!\!/ - m_N)N + (d\bar{\nu}_L\sigma_{\mu\nu}F^{\mu\nu}N + h.c). \quad (1)$$

Here $F^{\mu\nu}$ is the electromagnetic field strength tensor, and ν_L is a SM neutrino field. This is an effective Lagrangian that needs to be UV completed at energy scales not much larger than $\Lambda \sim d^{-1}$. We are going to stay on the effective field theory grounds, noting that since our results show the sensitivity to d to be much better than TeV^{-1}, the UV completion scale can be raised above the electroweak scale. For now, Eq. (1) is also applicable only at energies below the weak scale, as it does not respect the full SM gauge invariance. Indeed, $F^{\mu\nu}$ should be a part of the U(1) and/or SU(2) field strength, and the insertion of the Higgs field H is also required, so that $d \propto \langle H\rangle\Lambda^{-2}$. For most of our analyses we will be interested in values of m_N in the interval from 1 MeV to 100 GeV, and at relatively small energies, so that a treatment using Eq. (1) is indeed sufficient.

The main assumption made in Eq. (1) is the absence, or subdominance, of the mass mixing operator NLH. When the mass mixing operator is dominant, the production and decay of N particles is mostly governed by its interaction with the SM particles via weak bosons. The phenomenological consequences of these minimally coupled particles N is well understood. In contrast, if the leading order operator is suppressed, the dipole operator offers novel signatures and features in the production and decay of N, such as a much enhanced role of electromagnetic interactions in the production and decay of N. This case has so far being addressed only in a handful of works [4–9], and here we would like to present a comprehensive analysis of the dipole N portal, and derive constraints on d that result from a variety of different experiments,

* gmagill@perimeterinstitute.ca
† plestird@mcmaster.ca
‡ mpospelov@perimeterinstitute.ca
§ yt444@cornell.edu

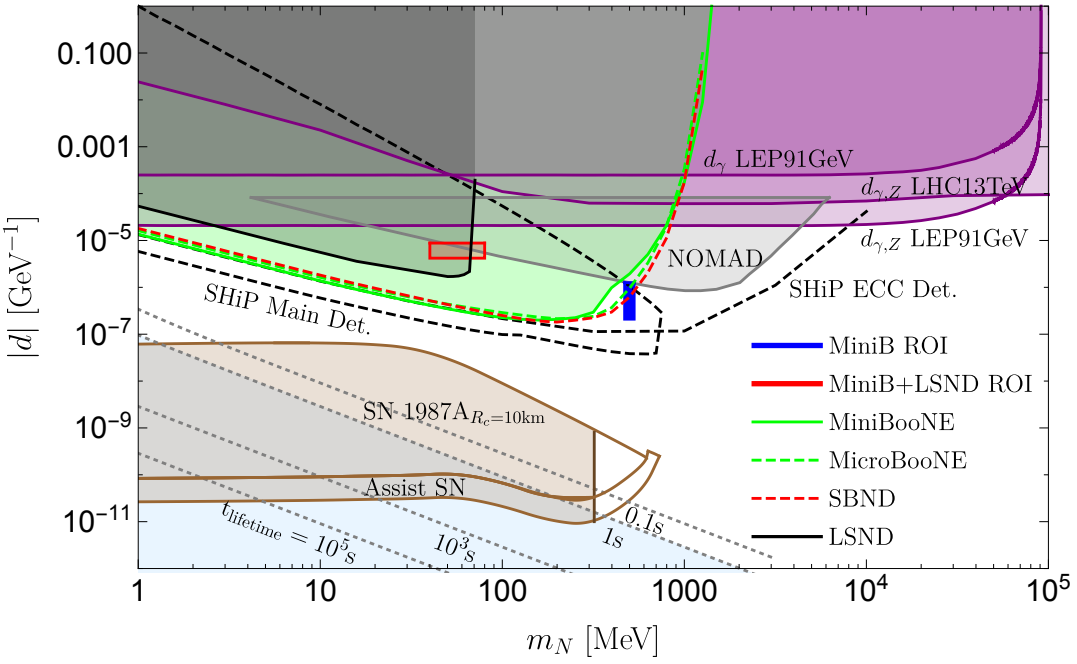

FIG. 1. Overview of projected sensitivities (95% CL) and constraints obtained from SHiP, LHC, LEP, Supernova 1987A and experiments at the Short-Baseline Neutrino facility at Fermilab. We also show previously calculated favored regions of interest (ROI) in parameter space for MiniBooNE and LSND, and constraints from NOMAD. Limits are shown for the dimension 5 (γ mediator) and dimension 6 ($\gamma + Z$ mediators) extensions. See Table II for an explanation of the labels. Each curve is discussed and presented in the paper.

both at high and medium energies.

Previously dipole interactions of neutrinos have been studied in several specific contexts (that we are aware of). If the SM neutrinos have a large flavor off-diagonal EM dipole moment, the interaction of solar and reactor neutrinos may get enhanced. This provides stringent limits on dipole moments of SM neutrinos [10]. Some theoretical and phenomenological aspects of the Dirac HNL dipole operator were discussed in Refs. [11, 12] (see also a more recent general discussion of dimension 5 effective operators in the neutrino sector [13]). A phenomenological sensitivity study of this magnetic dipole operator has been considered for IceCube [14]. There, owing to the large incoming SM neutrino energies, the signature of interest was a coincident double energy deposit from the DIS production of N, and its subsequent decay. Another prominent place where the transitional $\nu - N$ dipole appears is the literature on searches of sterile neutrino dark matter via a dipole-induced decay $N \to \nu\gamma$ ([15] and references therein). A more closely related case to the topic of our study has arisen as a consequence of trying to accommodate MiniBoone and LSND anomalies, that we would like to discuss now in more detail.

While there is an overall theoretical/experimental consistency for the three-neutrino oscillation picture, there are several experimental results that do not fit in. Two notable exceptions are the anomalies observed at the intensity frontier experiments LSND and MiniBooNE [16, 17]. In these experiments, an excess of low energy electron (anti-)neutrinos have been observed, the source of which is currently unknown. Conceivably, there are two possibilities: new physics or some unaccounted SM processes. Thus, for example, single photons produced via poorly understood SM neutrino interactions with nuclei [18] might lead to some partial explanation of the anomalies. (At the signal level, a single photon cannot be distinguished from charged-current quasi-elastic events by MiniBooNE's Cherenkov detector.)

The most popular proposal is the existence of a light ($m \sim$ eV) sterile neutrino ([19] and references therein), which mediates the anomalous oscillation required to explain the observed excess signal. A possibility of eV sterile neutrinos being at the origin of the MiniBooNE and LSND oscillation results is strongly challenged by cosmological data. Indeed, the required parameters for mass splitting and mixing angle will lead to a complete thermalization of a new sterile species via oscillation mechanism. This stands in sharp disagreement with cosmological data (in particular, cosmic microwave background (CMB), Big Bang Nucleosynthesis (BBN)

and late-time cosmology) that constrain not only the total number of thermally populated relativistic degrees of freedom in the early Universe, but also limits the total neutrino mass $\sum m_\nu \leq 0.17$ eV at 95%CL [20]. Consequently, a single eV sterile neutrino is not consistent with cosmology in the absence of new physics. At the very least, the minimal model would need to be modified to suppress the oscillations in the early Universe, which is usually achieved at the expense of significantly enlarging the sterile neutrino sector *e.g.* by new types of interactions with dark matter and/or baryons [21, 22]. Thus, the sterile neutrino solution to the MiniBooNE and LSND anomalies naturally leads to the idea of a *dark sector*, with new matter and interaction states.

An alternative attempt to accommodate the anomalies without using eV-scale sterile neutrinos requires some dark sector states comparable in mass to the lightest mesons. Thus, it has been noted that the presence of a new sub-GeV neutral fermion N may mimic the signals observed at MiniBooNE and LSND [4, 5]. The necessary ingredient of this proposal is a new fermionic state N in the 10-to-few-100 MeV mass range and the dipole coupling in Eq. (1). This coupling mediates a relatively prompt decay of N to a normal neutrino and a photon, a signature that can be confused with the "normal" electron or positron final state in charged current events [4, 5]. Whether this model can simultaneously account for both anomalies without running into problems with other constraints remain an open issue (see the discussions in Refs. [4–9]). At the same time the model has a clear advantage over the eV sterile neutrino model, as it creates no problems with cosmology, as N states will decay to the SM at early times before the neutrino decoupling.

Continuing investment in neutrino physics will eventually lead to better understanding of the origin of these two anomalies. The Short-Baseline Neutrino program (SBN) [23] is going to be instrumental in testing the MiniBooNE anomaly. The design consists of three Liquid Argon time projection chamber (LAr-TPC) detectors that overcome the difficulties present at MiniBooNE by providing excellent photon/electron discrimination. Furthermore, the SBN program will use a near detector (SBND) to control systematic errors related to the neutrino beam content. Being close to the proton target, SBND will see a much larger neutrino flux than the mid-range detectors and will allow a more accurate measurement of the neutrinos before oscillation. In addition, a further increase in sensitivity may result from a proposed new experiment at CERN, Search for Hidden Particles (SHIP) [3], that will be able to significantly advance the probes to N states, and should also test their dipole interactions. For an analysis of a more conventional CC-dominated model of HNLs in application to Fermilab experiments we refer the reader to a recent paper [24].

Motivated by the relative simplicity of the neutrino dipole portal model and its potential applicability to neutrino anomalies, it is very useful to have a comprehensive survey of the model over a large region of parameter space. We therefore consider the energy, intensity and astrophysics frontiers, where this portal can be probed. A plot summarizing our results is shown in Fig. 1, and the rest of the paper considers each probe individually. The existing constraints from previous dark matter experiments can be improved by the SBN and SHiP. From astrophysics, MeV HNLs could contribute to the supernova cooling, in particular that of Supernova 1987A (SN 1987A). This happens when the coupling d is large enough so that the star can produce N in sufficient quantity, but small enough so that N can escape and cool down the star without being significantly impeded. For lifetimes longer than $0.1\text{s} - 1\text{s}$, N is relevant for, and can modify predictions of, BBN. The late decays of HNLs would modify the proton to neutron ratio, and with some reasonable assumptions about the initial cosmological temperatures being high, this puts an upper bound on the lifetime of N. We find that there is significant overlap of this region with SN constraints. Lastly, for above GeV masses, we turn to particle colliders and recast existing searches from the LHC and LEP. Going to particle colliders allows us to probe simple completions of the model which preserve the $\text{SU}(2) \times \text{U}(1)$ structure of the SM. In these extended models, we have additional production channels stemming from Z and W bosons.

The paper is organized as follows. In Section 2, we provide more details on the model including the possible SM gauge invariant completions and the connections to neutrino masses. In section 3, 4 and 5, we consider the intensity, energy and astrophysical frontiers respectively. Finally, we conclude in section 6 with general remarks.

II. GENERIC FEATURES OF NEUTRINO DIPOLE PORTALS

A. Main qualitative features of dipole portal

The consequences of the dipole portal in Eq. (1) can be easily understood by considering the four vertex alignments presented in Fig. 2. The presence of an electromagnetic coupling to neutrinos allows for mesons to decay in two novel ways: Dalitz-like decays mediated by off-shell photons and neutrinoless weak decays with a single photon in the final state. In terms of producing N, incident neutrinos can upscatter via the dipole portal, which can be a more efficient production process than mass mixing mechanisms that have been traditionally considered. The decay of an HNL in our model will be dominated by single photon production, and for the values of d's we consider in this paper, will occur much more rapidly than in mechanisms that are mediated by the weak force. This single photon signature was identified in Refs. [4, 5] as a promising signal, however the production mechanisms

outlined above were not included.

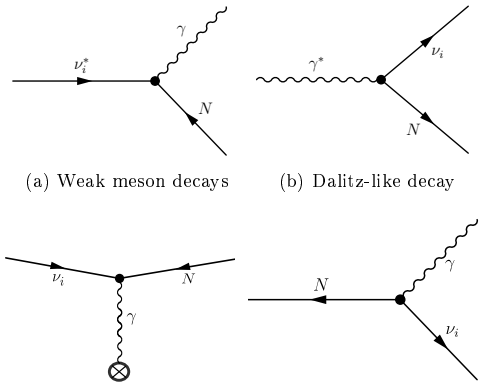

(a) Weak meson decays (b) Dalitz-like decay

(c) Primakoff upscattering (d) $N \to \gamma\nu$ (signal)

FIG. 2. Dipole portal processes for N: (a) Production of N from off-shell neutrinos arising from weak meson decays (*e.g.* from $\pi, K \to \mu^+\nu^*$); (b) Production of N from off-shell photons arising from Dalitz-like meson decays (*e.g.* from $\pi^0, \eta \to \gamma^*\gamma$); (c) Production of N from on-shell neutrinos via Primakoff-type upscattering (via photon exchange with the nucleus); (d) Decays of N to single photon final states (the main signal studied in this paper). Processes (a) and (b) are important for production of low mass N at neutrino experiments. Process (c) dominates production in supernovas at lower N masses, and at neutrino experiments. Process (d) is relevant for energy injection at BBN, for neutrino beam dump experiments, and controls the escape probabilities in supernova for large N masses.

We now focus our discussion to beam dump experiments. There, production of N will dominantly proceed via neutrino upscattering, wherein an incoming neutrino scatters via a photon to produce N. If the incoming neutrino scatters off the whole nucleus and the process happens coherently (i.e. $\sigma \propto Z^2$), we can get a crude estimate for the sensitivity one can achieve. In the limit of infinite mass of the nucleus, the problem reduces to the scattering in the external EM field $A_\mu = (A_0(\vec{q}), 0)$ created by the nucleus. Calculating the cross section to logarithmic accuracy for $-\frac{1}{R_{\rm nuc}^2} \leq t \leq -\frac{m_N^4}{4E_\nu^2}$, we find

$$\sigma^E_{\nu\to N} = 4\alpha Z^2 |d|^2 \times \log\left(\frac{4E_\nu^2}{m_N^4 R_{\rm nuc}^2}\right), \tag{2}$$

where we have retained the leading Z-enhanced contribution that corresponds to the interaction of νN dipole with the electric current created by the nucleus. For reference, we also include the expression for $\nu \to N$ upscattering due to the magnetic moment of the nucleus. Cutting off the coherent scattering at $|t| = \frac{1}{R_{\rm nuc}^2}$, one has

$$\sigma^M_{\nu\to N} = \frac{(I+1)}{3\pi I}\mu_A^2 |d|^2 \times \frac{1}{R_{\rm nuc}^2} < \sigma_E/Z^2, \tag{3}$$

where μ_A is the magnetic moment of the nucleus $\propto \alpha^{1/2} m_p^{-1}$. For a typical Z involved, the contribution of the nuclear magnetic moment to the coherent $\nu \to N$ upscattering can be safely neglected. More detailed and related discussions of dipole interactions of dark sector particles can be found in [25].

Therefore for masses $m_N = 50$ MeV, an incoming neutrino energy of 1 GeV and $R_{\rm nuc}^{-2} \sim 0.3$ GeV^2, we can expect a production cross section per nucleus of roughly

$$\sigma = 4.5 \times \left(\frac{Z}{18}\right)^2 \left(\frac{|d|}{10^{-6}\ \text{GeV}^{-1}}\right)^2 \times 10^{-38}\ \text{cm}^2. \tag{4}$$

It is worth noting that the $|d|^2$ scaling of the cross section makes it also a relatively mild, logarithmic function of energy, provided that N is kinematically accessible.

For characteristic values of d suggested in Eq. (4) and small masses, we can also expect N to be long-lived. This opens up the possibility of HNLs being produced outside of the detector. For example, they could be produced in the dirt or line of sight leading up to the detector, and/or via mesons from the protons-on-target via Dalitz-like decays. Meson production via the dipole portal is an interesting new production mechanism we will discuss, and from dimensional arguments it is clear that the scaling of the meson decay branching to N will occur via $Br_{M\to N} \propto d^2 m_M^2$, where m_M is the mass of the decaying meson.

The decay length associated with the $N \to \nu\gamma$ process is another very important quantity. Given a decay rate of

$$\Gamma_{N\to\nu\gamma} = \frac{|d|^2 m_N^3}{4\pi}, \tag{5}$$

and an HNL energy of $E_N = 1$ GeV $\gg m_N$, the decay length and lifetime of N scale as

$$\begin{aligned} t_{\rm dec} &= \tau\gamma = 1.3 \times 10^{-6}\text{s} \left(\frac{50\ \text{MeV}}{m_N}\right)^4 \left(\frac{10^{-6}\ \text{GeV}^{-1}}{|d|}\right)^2 \\ L_{\rm dec} &= c\tau\beta\gamma \approx 400\text{m} \left(\frac{50\ \text{MeV}}{m_N}\right)^4 \left(\frac{10^{-6}\ \text{GeV}^{-1}}{|d|}\right)^2. \end{aligned} \tag{6}$$

This turns out to be a very convenient length scale for beam dump experiments, if m_N and d have the fiducial values suggested above.

B. Dirac vs Majorana masses and gauge invariant completions

If N_D is a Dirac fermion, composed of two Weyl fields

$$N_D = \begin{pmatrix} N \\ N^{c\dagger} \end{pmatrix}, \tag{7}$$

one of which is completely decoupled from the SM, then the HNL is decoupled from the mechanism that generates active neutrino masses. Thus, we assume both the

absence of mass mixing between ν and N, and a vanishing Majorana mass for N. This choice is technically natural and can be achieved by—for example—assigning N the same lepton number as the SM leptons. If such a symmetry is not imposed, and a sizeable Majorana mass term, $\mathfrak{m}_N$, is present then the process shown in Fig. 3 can take place. Naive counting of divergences shows that the induced Majorana mass for the neutrinos, $\mathfrak{m}_\nu$ will scale as $\mathfrak{m}_\nu \sim d^2\Lambda^2\mathfrak{m}_N/16\pi^2$, where Λ is the cutoff scale associated with the UV completion of the model, which can be as high as d^{-1}. This contribution, despite all the uncertainties, will be much larger than the required mass scale for the neutrinos, unless N is Dirac, or quasi-Dirac with a small Majorana-type mass splitting satisfying $\mathfrak{m}_N \ll m_N$. Quasi-Dirac N would typically lead

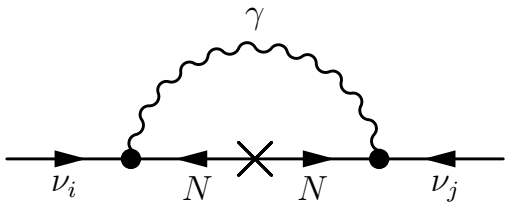

FIG. 3. Loop level contribution to the ν mass mixing matrix in the presence of a Majorana mass term for the heavy neutral lepton N. With only Dirac masses, such diagrams will not be generated.

to *larger* values of d than otherwise would be suggested by a simple application of the see-saw relation. Consider a model where the SM neutrinos couple to N via a mass mixing interaction of the form $m_{\nu N}\ \nu N$. This naturally generates dipole couplings between the SM neutrinos, sterile neutrino and the photon via a loop diagram. The dipole coupling generated is given in [26–28] as

$$\begin{aligned} d &= \frac{3m_{\nu N}}{32\pi^2}\frac{eG_F}{\sqrt{2}} \\ &= 1.2\times 10^{-9}\ \text{GeV}^{-1}\left(\frac{m_{\nu N}}{50\ \text{MeV}}\right). \end{aligned} \tag{8}$$

The strength of this radiatively generated dipole portal is dictated by the mass mixing with the active neutrinos, and therefore constrained by patterns of the neutrino mass matrices. In particular, in the case of a type-I see-saw mechanism with the Majorana mass of $m_N = 50$ MeV, observed neutrino masses would imply $m_{\nu N} \sim$ keV and consequently $d \sim 10^{-13}$ GeV^{-1}. We do not impose such a stringent constraint and consider d to be an independent parameter. In fact, the size of d can be much larger if the effective mixing angle between ν and N is much larger than the naive see-saw relation implies. This may happen, for example, within an inverse see-saw model [29, 30], where a mostly Dirac fermion N is supplemented with a small Majorana mass, so that the mass mixing parameter $m_{\nu N}$ is much larger than naively implied.

Above the electroweak scale an SU(2) × U(1) interpretation of d would require a Higgs insertion, so that the dipole iTnteraction is really a dimension 6 operator. Therefore, in the limit of large Λ the maximum expected d is

$$d_{\max} \sim \frac{e\text{v}}{\Lambda^2} \sim \frac{100\ \text{GeV}}{\Lambda^2} \tag{9}$$

where strong dynamics at the scale Λ is presumed, and v is the Higgs field vacuum expectation value. Otherwise, if the new sector is perturbative, we would expect a loop factor, and $d_{\max,\ \text{pert}} \sim$ GeV$/\Lambda^2$. To consider neutrino dipole couplings which respect the full gauge symmetries of the Standard Model, we write down the Lagrangian

$$\mathcal{L} \supset \bar{L}\left(d_{\mathcal{W}}\mathcal{W}^a_{\mu\nu}\tau^a + d_B B_{\mu\nu}\right)\tilde{H}\sigma_{\mu\nu}N_D + h.c. \tag{10}$$

where $\tilde{H} = i\sigma_2 H^*$ and $\tau^a = \sigma^a/2$. After spontaneous symmetry breaking of the Higgs, one obtains

$$\begin{aligned} \mathcal{L} \supset & d_W\left(\bar{\ell}_L W^-_{\mu\nu}\sigma^{\mu\nu}N_D\right) \\ & + \bar{\nu}_L[d_\gamma F_{\mu\nu} - d_Z Z_{\mu\nu}]\sigma^{\mu\nu}N_D + h.c. \end{aligned} \tag{11}$$

where $W^-_{\mu\nu} \equiv \partial_\mu W^-_\nu - \partial_\nu W^-_\mu$. The dipole couplings in the broken phase are related to those in the unbroken phase via

$$\begin{aligned} d_\gamma &= \frac{\text{v}}{\sqrt{2}}\left(d_B\cos\theta_w + \frac{d_{\mathcal{W}}}{2}\sin\theta_w\right) \\ d_W &= \frac{\text{v}}{\sqrt{2}}\frac{d_{\mathcal{W}}}{2}\times\sqrt{2} \\ d_Z &= \frac{\text{v}}{\sqrt{2}}\left(\frac{d_{\mathcal{W}}}{2}\cos\theta_w - d_B\sin\theta_w\right) \end{aligned} \tag{12}$$

where the additional factor of $\sqrt{2}$ in the expression for d_W is a consequence of the normalization of $W^- = (\mathcal{W}^1 + i\mathcal{W}^2)/\sqrt{2}$. Note that the three "dipole moments" in the broken phase d_γ, d_Z and d_W are determined by only two parameters in the unbroken phase $d_{\mathcal{W}}$ and d_B; they are linearly dependent. Notice that the normalization of the photon field strength term in Eq. (11) matches that of Eq. (1).

Although we have suppressed the relevant indices, the dipole coupling can be flavor dependent. Experiments at SBN will constrain d^e_B and d^μ_B. SHiP in addition will be sensitive to ν_τ, and thus an ideal setting to study all "dipole couplings". For both LHC and LEP, we turn on only the $d^\mu_{\gamma,B,W}$ coupling for simplicity. One can also turn on $d^e_{\gamma,B,W}$ and $d^\tau_{\gamma,B,W}$ that have an O(1) effect on the result.

Having established that a neutrino dipole portal is ultimately a dimension 6 operator, one might wonder if there are any non renormalizable SM only operators that are phenomenologically equivalent to our new physics signal. If so, one would need to perform a

global fit on the whole basis of Wilson coefficients instead of focusing on just one operator. The case of SM only operators after electroweak symmetry breaking is considered in Section IIIB. Ref. [31] on the other hand provides a classification of all dimension 5 and 6 SM only operators above the electroweak scale (i.e. invariant under SU(3) × SU(2) × U(1)). In order to replicate our signature, we need at least one photon, one neutrino and an additional gauge boson. If we assume that no particles except neutrinos escape detection, and furthermore that the interactions are $2 \to \mathrm{N}$, then none of the dimension 5 or 6 operators in Ref. [31] contribute to single photon processes at beam dump experiments, LEP or the LHC.

Lastly, we would like to comment on electric dipole operators and CP invariance. For the theories studied in this paper, the "electric" and "magnetic" connection to CP properties is no longer straightforward. We deal here with fields of certain chirality (SM neutrinos are left-handed, and interact with right-handed N), and therefore the operator that we wrote down in the Lagrangians is unique owing to the fact that γ_5 can always be re-absorbed in the right-handed projection operator of the HNL (i.e. $\gamma_5 * P_R = P_R$). In that sense, we have not reduced the number of possible operators; the electric form-factors and magnetic form-factors in our parametrization are inseparable. CP violation, on the other hand, can still be present due to a possible relative phase difference between the dipole coupling d and the mass term which we take to be real. The processes we study are only sensitive to $|d|$, and not its phase.

III. INTENSITY FRONTIER

We consider probing HNLs at beam dump experiments and our analysis focuses on neutrino experiments hosted at CERN, Los Alamos and Fermilab. Fermilab is building a substantial Short-Baseline Neutrino oscillation program [23] that among other physics goals will settle the question of sterile neutrinos at $\Delta m^2 \sim 1$ eV^2. It will consist of 3 LAr-TPC detectors called SBND, MicroBooNE and ICARUS, which will be spread out over a 600m range from the proton target. The SBN program is designed to achieve a 5σ sensitivity in the parameter space of (3+1) sterile neutrino models consistent with LSND at 99%CL. These detectors can resolve photons from electrons with a 94% photon rejection rate.

At CERN, we will be interested in the past experiment NOMAD and future proposal SHiP. The proposed SHiP experiment is unique among beam dump experiments in that it features very large neutrino energies and a sizeable flux of electron, muon and tau neutrinos. Furthermore, the use of lead inside the neutrino detector, $Z = 82$, will provide an ideal setting to take advantage of coherent production, which scales as Z^2. At Los Alamos, we consider the LSND experiment which will prove to be useful at low HNL masses. In what follows, we discuss the various production mechanisms at beam dumps, the main backgrounds involved in the search, and our results.

A. Production mechanisms

At neutrino beam dump experiments, HNL production can happen in three principle ways. The first—and most familiar— mechanism is mass mixing, however this is subdominant in our analysis by assumption. The two dominant production mechanisms are therefore meson decays and Primakoff upscattering, both of which are explained in greater detail below. In principle DIS production via Drell-Yan like processes is also possible, but we found this to be subdominant.

1. *Primakoff upscattering*

Neutrino upscattering is the dominant production mechanism for N across a wide range of masses for the experiments we consider. It happens when an incoming neutrino interacts with matter and upscatters into a long-lived HNL state N. The HNL subsequently decays into a neutrino and a photon; an explicit example is provided in Fig. 4. This process can either happen inside the fiducial

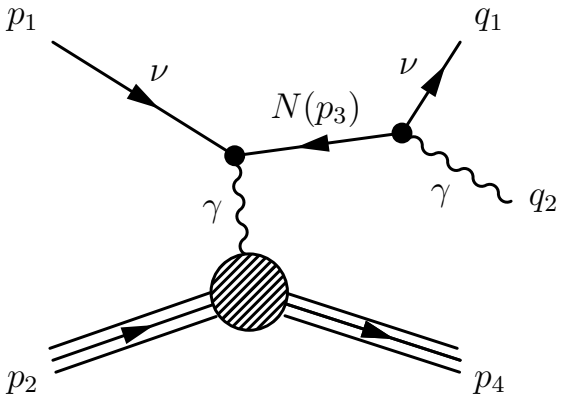

FIG. 4. Tree level neutrino scattering process with a final state photon, arising from dipole portal to HNL. We work in the narrow width approximation, and assume the above diagram factorizes.

volume of the detector or in the line of sight separating the proton target from the detector. In all our results, we employ the narrow width approximation, since N is usually produced on-shell and travels some distance before decaying. Having an HNL lifetime and energy consistent with the necessary flight distance is enforced by

$$P_{\mathrm{dec}}(L_1, L_2) = \exp[-L_1/L_{\mathrm{dec}}] - \exp[-L_2/L_{\mathrm{dec}}]. \quad (13)$$

In Appendix A 1, we present the details of how the cross section is obtained for coherent and diffractive scattering. We apply the cuts described in Appendix C to ensure proper kinematics of the photon. There, it is also fully described how the region of integration of t is determined. Once we have obtained the cross section, cuts, photon detection efficiency and luminosity, we can set limits following the discussion in Appendix B.

2. *Meson decays*

At low mass, HNLs are long lived and represent a kinematically allowed decay channel for light mesons. Unlike mass mixing induced decays, the dipole portal allows for electromagnetically mediated Dalitz-like pathways in addition to weak decays mediated by an off-shell neutrino. The qualitative features that can allow for significant production of HNLs are

(i) High meson multiplicity per proton (*e.g.* pions).

(ii) BR(Meson $\to X+\gamma$) $= \mathcal{O}(1)$ (*e.g.* π^0, η) or BR(Meson $\to X+\nu$) $= \mathcal{O}(1)$ (*e.g.* $\pi^\pm$, K).

In terms of meson production at the experiments we consider, the largest difference between them is that immediately following the proton target, SBN features a 50m meson decay chamber, whereas SHiP has a hadron stopper. This divides our discussion into prompt ($\tau^{\rm rest} \lesssim 10^{-12} s$) and long-lived ($\tau^{\rm rest} \gtrsim 10^{-12} s$) mesons. Only the former will contribute to HNL production at SHiP, whereas both will be relevant at SBN due to its long decay chamber. To obtain rates, we calculate the differential cross section of HNL production from mesons in the meson rest frame, which we combine with the meson fluxes in the lab frame. The details of these calculations are outlined in Appendix A 2.

The species we have included in our analysis are shown in Table I, from which it is clear that the prompt mesons are π^0 and η. For both of these, the dominant channel for HNL production is

$$\pi^0, \eta \to \gamma(\gamma^* \to \nu_a N) \tag{14}$$

We immediately see that these radiative Dalitz-like decays will be useful for improving the sensitivity to d_e and d_τ flavored couplings, since the process in Eq. (14) is universal in the flavor a. By contrast, neutrino upscattering at beam dump experiments is limited by smaller incident fluxes of ν_e and ν_τ neutrinos, as compared to ν_μ neutrinos. The long-lived mesons we consider are $\pi^\pm$ and $K^\pm$. They can produce HNLs via an off-shell neutrino decay

$$\pi^\pm, K^\pm \to \mu^\pm \left(\overset{(-)}{\nu}_\mu{}^* \to \gamma \overset{(-)}{N} \right). \tag{15}$$

When considering decays to electron flavor, such as $K, \pi \to e\nu_e$, one typically expects a chiral suppression of $\mathcal{O}(m_e^2/m_\mu^2)$ in the branching ratio relative to the muon channel. While we concentrate on Eq. (15) for muon flavors at SBN, we note that $K, \pi \to eN\gamma$ will avoid chiral suppression due to the chirality-flipping nature of the dipole portal. The K^+ states, whose rates are about a tenth of those of pions, are important because they allow production of heavier HNLs.

Meson Species	Multiplicity per POT	$\langle p \rangle$ [GeV/c]	$\langle \theta \rangle$ [mrad]	$\langle \tau \rangle$ [sec]
π^-	0.9004	0.83	527	$2.6 \cdot 10^{-8}$
π^+	0.9784	1.07	423	$2.6 \cdot 10^{-8}$
π^0	0.9098	0.89	483	$8.4 \cdot 10^{-17}$
K^+	0.0689	1.33	410	$1.2 \cdot 10^{-8}$
K^-	0.0024	1.29	409	$1.2 \cdot 10^{-8}$
η	0.0295	1.35	403	$5.0 \cdot 10^{-19}$

TABLE I. Meson multiplicities, average momentum and average angle at the SBN facility. Pions are assumed to follow a Sanford-Wang distribution, while kaons and etas are calculated based on the Feynman Scaling distribution.

To get a handle on which mesons are expected to contribute most, we calculated the average multiplicities of each meson per proton on target at SBN. Our results are shown in Table I. The π^- multiplicity has been calibrated to match that of Table X in [32], and we find very good agreement for the other meson multiplicities. No distribution parameters for K^- and η were available, and so we rescaled those of K^+ to match expectations. Both K^- and η contributions are very small, so the discrepancy in average momentum and angle as compared with Table X has a negligible effect on our results. We conclude that pions will be the most important mesons for sourcing low mass HNL particles.

B. Backgrounds

The main backgrounds for HNLs will be single photon signatures, arising from mis-reconstructed π^0 or radiative resonance decays such as $\Delta \to N\gamma$. At SHiP, there is not much publicly available information, and therefore we consider various benchmark estimates for these backgrounds. We guide our estimate by considering the observed single photon backgrounds at NOMAD, rescaled to account for differences in the target mass and number of protons on target.

On the other hand, the SBN collaboration has estimated the number of single photon events that can fake a ν_e CC signature in each of its detectors. We can estimate the total single photon background by taking this number and dividing it by 6% to factor out the photon rejection rate. We then impose a 200 MeV threshold in our

results since the single photon backgrounds grow with decreasing energy. To account for signal photons that may have been lost, we apply a 20% signal efficiency cut.

The backgrounds at LSND are similar in spirit to those at MiniBooNE, in that electron-like events arise from both electron and photon sources. In order to obtain constraints, we base our analysis at LSND on an electron-neutrino elastic scattering search [33], respecting the fiducial geometry and energy cuts described in that paper. Substantial constraints on various weakly interacting light particles, including scalars, dark photons, and fermionic dark matter, can be placed based on this measurement, with these light particles being produced on-shell and alter the neutrino-electron elastic scattering signature through their decays or scattering with the electrons, similar to the procedure discussed in this work [34–36]. Examining Fig. 1 and 10 of [33], we note that the incident neutrino flux favors energy values between 30-50 MeV, whereas the collected electron-like sample peaks at energies around 22 MeV. Single photons from HNL decays on the other hand tend to be much harder and closer in energy to their parent SM neutrino. We therefore explore two different recast strategies. In the first case, we impose a lower threshold on the incident neutrino energy of 18 MeV. This corresponds to the full dataset collected by LSND, comprising of roughly 300 predicted background and data events. In the second strategy, we impose a lower energy cut of 40 MeV in an attempt to better discriminate our new physics signal from SM backgrounds. This cut amounts to keeping roughly 27 predicted background and data events. We find that the latter strategy provides slightly better sensitivities to HNLs, and these are the LSND results that feature in all of our plots.

Lastly, diagrams containing loops of charged leptons and either a W or Z boson, can induce an effective $\gamma\gamma\nu\nu$ vertex in the SM and provide a potential source of single photon backgrounds. We have explicitly estimated the size of this background in Appendix A 3 and it is many orders of magnitude lower than the HNL production cross section estimated in the previous section, and can therefore safely be ignored.

C. Experimental results and prospects

In what follows we describe and summarize the implications of existing measurements at LSND and MiniBooNE. We also comment on the projected reach of ongoing and future experiments such as MicroBooNE and SHiP.

1. *LSND*

The LSND oscillation anomaly, which consists of an excess of $\bar{\nu}_\mu \to \bar{\nu}_e$ events [16], has historically motivated interest in sterile neutrinos. While common interpretations of the excess typically involve very light sterile states, more recently it has been proposed that a dipole portal coupled with HNLs with $m_N \approx 50$ MeV could explain the excess [4, 5, 7]. It is therefore of great importance to consider the observations at LSND and their implications for dipole portals to HNLs.

The setup at LSND involves a neutrino flux coming primarily from μ^+ and π^+ decays at rest [33]. Consequently the dominant production channel of HNLs is through neutrino upscattering. In modelling the production of HNLs at LSND we include Primakoff upscattering of neutrinos, as well as decays in flight for π^0, decays at rest for μ^+ and decays both at rest and in flight for π^+. We account for the change in LSND's source of neutrinos, LAMPF, and include two years of data assuming a water based target and three years of operation using a high-Z target (mostly tungsten) [33]. For our purposes the primary effect of the target material is to modify the incident flux of neutrinos, and mesons.

The decays in flight of π^+ and π^0 are modelled assuming a Burman-Smith distribution with appropriate parameters for both water and tungsten [37, 38]. Additionally, the decay at rest of μ^+ and π^+ contribute to the production of HNLs. The decay mode of interest for π^0 is a Dalitz-like decay, while for μ^+ and π^+ an off-shell neutrino mediates the production of HNLs. This off-shell neutrino can be either $\overline{\nu}_\mu$, or ν_e and we include both of these processes in our analysis. Summing all of these processes, and appropriately boosting the HNLs from decays in flight, leads to an incident flux of HNLs which may enter the detector and decay leaving a single photon signature.

On top of a flux of HNLs due to pion and muon decays, Primakoff upscattering of neutrinos in transit on their way to the detector can provide an additional source of HNLs. Alternatively, upscattering can occur within the detector itself. These processes have to be considered separately since much longer decays are possible in the case of the former, the target material upon which the neutrino upscatters is different, and angular cuts will be dictated by the different geometries.

When upscattering in transit to the detector, the medium of interest is the dirt—and other terrestrial material—along the line of sight between the source and the detector. In our analysis this is modelled as SiO_2 and we include both coherent and diffractive scattering. The produced HNL must be directed in a range of solid angle so as to guarantee that it passes through the detector. The range of angles for which this occurs is different depending on how far away the HNL is produced from the detector. To account for this effect, we analyze ten evenly spaced points between the source and the detector. At each of these points, given a flux of neutrinos, we calculate the number of HNLs that would both be produced *and* enter the fiducial volume of the detector. The LSND detector is off-axis from the neutrino source, and is roughly cylindrical in shape, and so we define the angular cuts such that the HNL would pass

through the bottom-near and top-far corners (relative to the neutrino source) of the detector; the angular cuts are implemented as described in Appendix C and account for fiducial cuts at the bottom of the detector. In addition to passing through the detector, the HNL's subsequent decay must occur within the fiducial volume for a signal to be observed. We account for this effect by including the probability that the HNL decays in the fiducial volume Eq. (13). Angular cuts within the detector are, as before, described in Appendix C.

It is also possible that upscattering occurs within the fiducial volume of the detector. For LSND this implies a target composed of CH_2 (mineral oil) for the incident neutrinos, and implies furthermore that neutrinos can be produced and subsequently decay along the entire line of sight. We account for this effect at leading order in the limit of $L_{\rm dec} \gg L_{\rm fid}$, which is the relevant regime when considering the minimal bound on the dipole-coupling of the HNL. We restrict the production of HNLs to the forward pointing hemisphere (*i.e.* an angular cut of $\theta \leq \pi/2$), due to experimental cuts. Additionally, we only include the effects of coherent scattering due to the presence of a hadronic veto within the detector.

2. *Fermilab's SBN program*

At Fermilab, we are interested in the past experiment MiniBooNE, as well as ongoing experiments involving the SBND and MicroBooNE detectors. At MiniBooNE, we consider the existing search for $\nu_\mu \to \nu_e$ quasi-elastic scattering events [17]. When limited to reconstructed neutrino energies of $475 < E_\nu^{QE} < 1250$ MeV, they find very good agreement between background and data. However, for energies between 300 and 475 MeV, MiniBooNE sees a persistent excess. MiniBooNE, being an oil based Cherenkov detector, cannot distinguish electrons from photons. A possible explanation for the excess [18] is from the $\Delta \to N\gamma$ process faking a ν_e signal. A direct chiral perturbation theory calculation finds these rates to be twice as big as data driven estimates from MiniBooNE.

The more exotic interpretations of the MiniBooNE and LSND anomalies [4, 5] involve additional single photons from new physics coming from an HNL model with a large dipole coupling d and an active neutrino mass mixing term in the range $|U_{N\nu}|^2 \simeq 10^{-3} - 10^{-2}$. In that case, production of HNL arises from neutral current ν scattering that leads to the production of HNL. In Fig. 5, we revisit the constraints from MiniBooNE by considering both production and decay stemming only from the dipole portal. Since it is difficult to reconstruct HNL energies (due to energy being carried away by outgoing neutrinos), we take an inclusive approach and sum over all the backgrounds and data bins. We calculate the allowed 95%CL HNL limits following the procedure in Appendix B for three different assumptions, which we denote by Bkg 1, 2 and 3 in Fig. 5. Firstly, we use the data and backgrounds as given in [17]. Secondly, we repeat the analysis after including the additional sources of backgrounds identified in [18]. And lastly, we compute constraints taking into account only the $E_\nu > 470$ MeV region. Based on [18], we assume a 25% photon identification efficiency to account for resolution and smearing effects. The photon energy detection threshold is 140 MeV. Comparing our results to [5] where dipole portal production mechanisms are ignored, we see that around 50 MeV masses production from dipole portal is actually dominant. An explicit calculation reveals that for the best fit parameters in [5], the dipole production cross section is roughly 20 times larger than production from mixing, and so this explanation appears to be excluded. This point is discussed in [7], and in the same work, the authors attempt to accommodate the constraint from the muon capture with photon emission at TRIUMF [6, 39] by introducing an additional heavy neutrino $\nu_{h'}$. In this way N can decay to $N \to \nu_{h'}\gamma$ as a main decay channel, and the branching ratio to ν_μ can be adjusted to accommodate the LSND/MiniBooNE anomalies while evading muon capture bounds. This same model was recently considered in the context of coherent and diffractive scattering at both MiniBooNE and MicroBooNE [40]. In contrast, we make no attempt to go beyond the minimal dipole coupling and we therefore exclude the favored regions of [5].

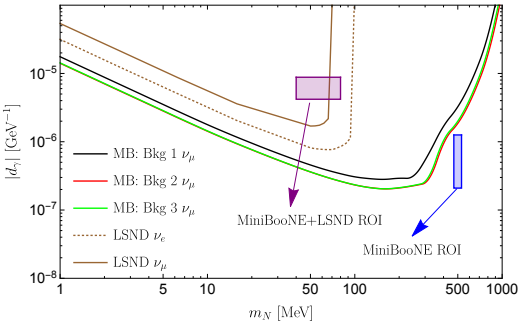

FIG. 5. 95% CL limits for HNL particles using MiniBooNE and LSND ν_eCC measurements. In light of the experimental anomaly, background option 1 (Bkg 1) uses the data and backgrounds as is, option 2 includes an alternative stronger $\Delta \to \gamma N$ background estimate [18], and option 3 includes only neutrino energies in the anomaly-free region ($E_{\nu_e} > 470$ MeV). We also overlay regions of interest (ROI) from the MiniBooNE and LSND anomalies (see text).

For 500 MeV HNL masses explaining MiniBooNE data, we find that production from mixing dominates. Therefore in order to obtain stronger dipole-only constraints, we turn to ongoing and future experiments. Our results for SBND and MicroBooNE are shown in Fig. 6. They assume 6.6×10^{20} POT of data in SBND and 13.2×10^{20} POT of data in MicroBooNE. As we see, after only 3 years of data taking, they can start cutting

into favored parameter space, provided photon data is collected in this duration.

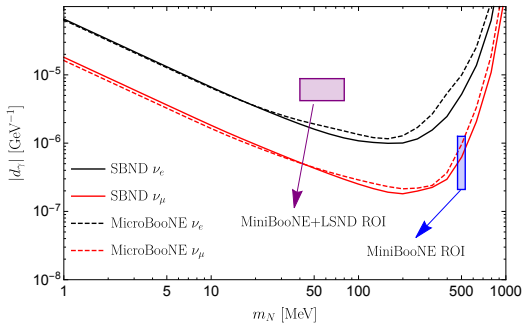

FIG. 6. Projected 95% CL sensitivities at Fermilab's upcoming Short-Baseline Neutrino program [23]. Results for electron (black) and muon (red) dipole couplings are shown for the SBND near detector (solid) and the MicroBooNE middle detector (dotted). Backgrounds are calculated based on expected lifetime single photons (see text).

3. *SHiP and NOMAD*

At CERN, we will consider the NOMAD experiment, which ran from 1995-1998 [41–43], and the proposed SHiP experiment. Both of these experiments are based on CERN's Super Proton Synchrotron, and consequently have neutrino fluxes extending to larger energies as compared to Fermilab. NOMAD has already performed a search for single photon production. Using this data corresponding to 1.45×10^{18} POT, Monte Carlo simulations of HNL signals (with no mass mixing) where performed [44, 45] to simulate the Primakoff process $\nu_\mu Z \to NZ$. The signature of interest was an isolated electromagnetic shower corresponding to a single photon with energy distributed from 0 to E_ν, with $E_{\nu_\mu}/2$ as an average. The backgrounds, estimated to be roughly 10 events, come mainly from π^0 production, as well as ν_e CC interactions. The full results[1] from their simulation are shown in Fig. 7.

CERN has also proposed a future high energy facility called SHiP [46]. If indeed funded and built, it would provide some of the strongest probes of heavy neutral leptons to date [3]. At SHiP, neutrinos are produced by 400 GeV protons impinging on a molybdenum and tungsten target. A hadron stopper immediately after the target allows only prompt meson decays, and a magnetized iron shield deflects muons. Following this is an emulsion cloud chamber near detector (which we will refer to as "ECC detector") containing lead bricks, a vacuum decay chamber followed by the main detector (which we will refer to as "main detector"). The length of the whole experiment would be on the order of 100m. It is advantageous to consider HNL production from prompt mesons, the line of sight, and lead bricks in order to maximize our sensitivity to a large range of HNL lifetimes. We apply a photon detection efficiency of 80% and an energy threshold of 0.1 GeV.

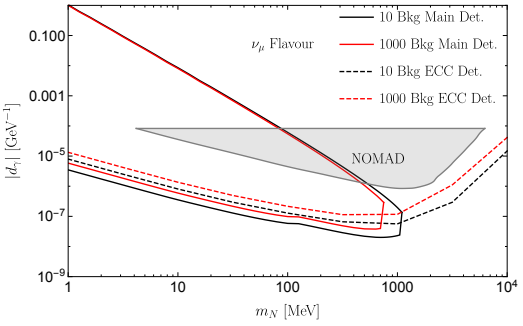

FIG. 7. Projected 95% CL sensitivities at SHiP for muon neutrino dipole moments. Solid (dotted) lines indicate the main (ECC) detector, and black (red) lines represent 10 (1000) background events during the lifetime of the experiment. We also overlay existing constraints [44, 45] from NOMAD.

A unique feature of SHiP is that it is expected to have a sizeable flux of ν_e and ν_τ neutrinos. Therefore, we can interpret the results of the single photon search as constraints on d_γ^f, for a given flavor f. Recall that flavor indices in Eqs. (1) and (10) are suppressed and a priori general. The projected sensitivities achievable at SHiP are shown in Fig. 7 for muon flavors assuming 2 different benchmark choices for the number of background events (10 and 1000 background events). In Fig. 8, we show the sensitivity for electron and tau dipole moments assuming 100 background events. At SHiP, single photon rates have not yet been studied. We can obtain a naive estimate by comparing to NOMAD, which had about 10 background events with 100 times less protons-on-target than SHiP. Therefore, with higher luminosities coupled to improved detector capabilities, it is reasonable to estimate around 100-1000 background events in the SHiP ECC detector. This detector will probably have more background events than the main detector, since the latter is surrounded by veto structures designed to reduce backgrounds as much as possible. For the SHiP curves appearing in Fig. 1, we assume 1000 backgrounds events in both detectors in order to provide a conservative estimate.

[1] The dipole coupling in their paper, ($\mu_{\rm trans}$), differs from ours by a factor of $\sqrt{2}$ ($\mu_{\rm trans} = \sqrt{2}d$).

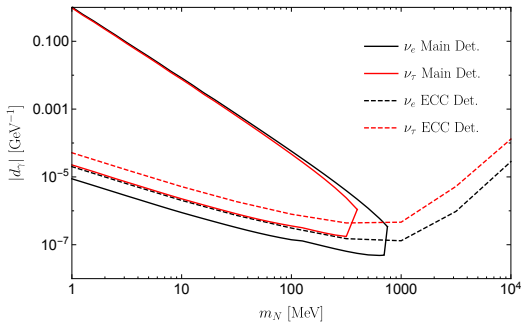

FIG. 8. Projected 95% CL sensitivities at SHiP for electron (black curve) and tau (red curve) neutrino dipole moments. Solid (dotted) lines indicate the main (ECC) detector. In this plot, we assume 100 background events.

IV. ENERGY FRONTIER

A. Production mechanisms

Beam dump experiments feature very large luminosities, however, the masses of N which are accessible are limited by the incoming neutrino energy spectrum, typically peaked around 1 GeV, or between $10-20$ GeV in the case of SHiP. In contrast, particle colliders can probe much larger masses at the expense of smaller luminosities [11]. Additionally, since dipole operators must couple to either $B_{\mu\nu}$ or $W_{\mu\nu}$ above the electroweak scale there is the added possibility of on-shell production of the Z and W mediators. The HNL couplings appearing in all of the high energy plots for LEP and the LHC are defined as follows. We take the relations in Eq. (12) and rescale $d_{B,\mathcal{W}} \equiv \sqrt{2}\,\overline{d_{B,\mathcal{W}}}/(\mathrm{v}\cos\theta_w)$ to obtain

$$\begin{aligned} d_\gamma &= \overline{d_B} + \frac{\tan\theta_w}{2}\overline{d_{\mathcal{W}}} \\ d_W &= \frac{\overline{d_{\mathcal{W}}}}{\cos\theta\sqrt{2}} \\ d_Z &= \frac{\overline{d_{\mathcal{W}}}}{2} - \tan\theta_w\overline{d_B}. \end{aligned} \tag{16}$$

Table II illustrates the assumptions made in each of the exclusion curves for LEP and the LHC.

We now discuss the mechanisms for producing HNLs at LEP and the LHC, and then discuss the details of the analyses and our results.

1. LEP

At LEP, production will proceed via $e^+e^- \to (N \to \gamma\nu)\bar{\nu} + h.c.$. The signature to look for is thus a single photon final state with missing energy. This channel can proceed via either Z or γ mediators depending on the dipole coupling in the unbroken phase (see Eqs. (11) and (16)). Therefore the total production cross section at $s = m_Z^2$ for $e^-e^+ \to N\bar{\nu}$ integrated over all angles is

$$\begin{aligned} \sigma_{N\nu} = &\frac{\alpha\left|\overline{d_B}\right|^2\left(m_N^2 - m_Z^2\right)^2\left(2m_N^2 + m_Z^2\right)}{6\cos^2\theta_w\sin^2\theta_w m_Z^6\Gamma_Z^2}\times \\ &\left(\tan^2\theta_w m_Z^2\left(C_A^2 + C_V^2\right) + 4\cos^2\theta_w\sin^2\theta_w\Gamma_Z^2\right), \end{aligned} \tag{17}$$

where we treat the electron as massless and assume that $\overline{d_{\mathcal{W}}} = 0$. The axial and vector couplings are defined as $C_A = -1/2$ and $C_V = -1/2 + 2\sin\theta_W$. In practice, we apply the experimental angular photon and energy cuts described in Section IV B and Appendix C and do not make approximations on the masses of electrons.

Exp.	Plot Label	Assumptions	Probed d
LEP	d_γ	$\overline{d_{\mathcal{W}}} = 0,\ d_Z = 0$	$\overline{d_B}$
	$d_{\gamma,Z}$	$\overline{d_{\mathcal{W}}} = 0$	$\overline{d_B}$
LHC	$d_{\gamma,Z}$	$\overline{d_{\mathcal{W}}} = 0$	$\overline{d_B}$
	$d^a_{\gamma,W}$	$d_\gamma = a \times \overline{d_{\mathcal{W}}}$	$\overline{d_{\mathcal{W}}}$

TABLE II. Assumptions and conventions used in obtaining constraints at LEP and the LHC for the minimal HNL models and the HNL extensions respecting the $\mathrm{SU}(2)\times\mathrm{U}(1)$ symmetry of the Standard Model.

2. LHC

At the LHC, there are two main production channels we can consider. The first channel is analogous to LEP, and consists of oppositely charged quarks and antiquarks interacting via an s-channel photon or Z boson: $q_i\bar{q}_i \to (N \to \gamma\nu)\bar{\nu} + h.c.$. This gives the same signature as LEP, up to subtleties that will be discussed in Section IV B. In addition to neutral currents, the LHC provides us with the opportunity to study interactions proceeding via charged currents. The charged current couplings appeared as one of two possible couplings above the electroweak scale in Eq. (11), and leads to a final state consisting of a single photon, charged lepton and missing energy—for example: $u_i\bar{d}_j \to (N \to \gamma\nu)\ell^+$. For the LHC, the rate of production of HNLs is calculated using MadGraph5_aMC@NLO v2.5.5 [47], making use of FeynRules2.3 [48, 49] to load our implementation of the HNL model.

B. Experimental results and prospects

1. LEP

There have been many analyses dedicated to the $\gamma + E_{\mathrm{miss}}$ final state [50–53]. We choose to focus on the re-

sults of LEP1, which ran at a center of mass (COM) energy corresponding to the Z pole and accumulated about 200pb^{-1} of data, and LEP161 which ran at a COM energy of 161 GeV and accumulated 25pb^{-1}[54]. Using partial luminosity and combining many analyses, LEP1 was able to set an upper bound of 0.1pb on the cross section of new physics contributing to the $\gamma + E_{\rm miss}$ final state, within the angular acceptance range of $|\cos\theta_\gamma| \leq 0.7$ and requiring the outgoing photon to have a minimal energy of 0.7 GeV. We also enforce that the HNL decays within 1m of the interaction point using Eq. (13). To set constraints using LEP data that extend to slightly larger HNL masses, we point out that LEP's 161 GeV run also set an upper bound of 1pb on the single photon cross section from new physics.

2. LHC

To probe the coupling d_Z, we recast a recent dark matter search at $\sqrt{s} = 13$ TeV by ATLAS [55] involving final states containing at least one photon with $E_T^\gamma > 150$ GeV, missing energy greater than 150 GeV, and 0 or 1 jets. Events in our MadGraph simulation were generated with 0 or 1 photon, and no jets. Owing to the systematic uncertainties in the modelling of initial state radiation, only background predictions with 1 jets are shown in the ATLAS paper. We use a data-driven method to estimate the background events with 0 jets by looking at the ratio of data events reported to contain either 0 or 1 jet. Following this, we see a deficit of data events in both the 0 and 1 jet channels as compared to the background predictions, which will motivate us to adopt the CL_s method for estimating the sensitivity at the LHC, which we describe in Appendix B. The dominant background for this search was the irreducible $Z(\to \nu\nu)\gamma$ process, followed by $W(\to \ell\nu)\gamma$ in which the final state lepton was not detected. In addition to all the cuts described in the paper, we also impose a probability function requiring the HNL to decay before the closest distance to the ECAL barrel, namely $r = 1.5$m from the beamline. We take the photon ID efficiency to be 92% [56].

The LHC also provides us with the opportunity to probe the charged current HNL extension. We make use of a $\sqrt{s} = 8$ TeV CMS search for supersymmetric models with gauge-mediated breaking [57]. In its analysis, the collaboration searched for 1 electron/muon with transverse momentum greater than 25 GeV, 1 or more photons, and missing energy greater than 120 GeV. The dominant backgrounds in this search were misidentified photons, misidentified leptons, and electroweak backgrounds. In the case of CMS, the transverse distance from the beamline to the ECAL barrel is 1.29m, and the detection efficiency for electrons and muons are 80% and 90% respectively. There are no requirements on the number of jets, however they show results consistent with low jet activity by requiring that the scalar p_T sum of jets (H_T) be smaller than 100 GeV. In our event generation, we do not consider associated jet production, which provides us with a conservative estimate. We simulate production of N and ℓ from a W boson via the $d^a_{\gamma W}$ coupling, and decays of N to a neutrino and photon via the d_γ coupling. We do this for various relative magnitudes between $d^a_{\gamma W}$ and d_γ.

In both the CMS and ATLAS searches, results are shown in terms of several signal regions defined by an additional requirement on the missing energy. We cycle through each of these signal regions and independently calculate the sensitivity in order to find the most constraining missing energy requirement. We now briefly comment on ways in which one could extend the reach of this analysis. Access to longer HNL lifetimes could be achieved by using the location of the photons hitting the ECAL barrel and endcaps, and statistically mapping these back to the original direction of the HNL. Then, on an event-by-event basis, we could select different maximal distances in the probability of decay cut. Currently, we only used the distance of closest approach between the IP and the ECAL barrel. An additional possibility is to allow the HNL to decay somewhere inside of the ECAL as opposed to before reaching the surface. To avoid potential difficulties with triggering however, this might have to be done in association with jets or leptons. Lastly, tau flavored couplings could be explicitly probed in the $\ell+\gamma+\not{E}_T$ analysis by tagging tau leptons. This would be a nice complement to neutrino beam dump experiments, whose characteristic energies and neutrino flavors often prohibit tau production. We do not include tau leptons in our simulations.

3. Results

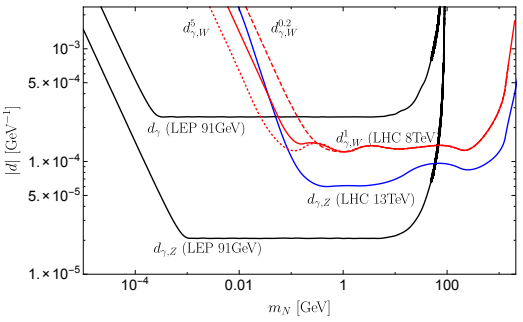

FIG. 9. 95% CL sensitivities at LHC and LEP. Limits are shown for the dimension 5 (γ mediator) and dimension 6 (γ, Z and $W^\pm$ mediators) extensions. For the LHC 8 TeV results involving a photon and charged lepton final state, we consider various relations between the production ($d^a_{\gamma W}$) and decay (d_γ) couplings. See Table II for an explanation of the plot labels.

The compilation of the high energy limits on the dipole couplings is presented in Fig. 9. All constraints have a characteristic "U" shape. The right boundary of the excluded region is controlled by the kinematic reach, and in the case of the LHC extends beyond a TeV. The left boundary (small m_N) is controlled by the lifetime of N, as smaller m_N leads to the longer lifetime of N and the loss of the γ signal in the detector. The bottom part of the constraints is controlled by the rates and backgrounds, and is approximately independent on m_N as in this region the production cross section is m_N independent, and its decay is relatively prompt. It is interesting that below $m_Z/2$ the LEP experiments are still capable of providing better sensitivity to the neutrino dipole portal.

V. COSMOLOGY AND ASTROPHYSICS

A. Big Bang Nucleosynthesis

Cosmology provides a very sharp tool in limiting the coupling constants of metastable heavy particles. In particular, consistency of BBN-predicted ^{4}He and deuterium yields with observations shows that the Universe was dominated by electrons, photons and SM neutrinos at very early epochs with temperature $T \sim 1$ MeV. Any massive relic surviving in large abundances down to these temperatures, or conversely having a lifetime in excess of 0.1 seconds, will distort this balance, and contribute to the Hubble rate during the proton-neutron freeze-out. Since most of the neutrons end up in ^{4}He, this possibility constrains the lifetime of heavy metastable relics *if* they are populated to large thermal abundances.

Therefore, we are led to investigate the mechanisms that populate HNLs in the early Universe. The analysis of the conventional mass-mixed case in its impact on BBN was performed in Ref. [58], and the mechanisms for thermal population of HNLs through neutrino oscillations is quite established [59]. Here we notice that the processes that populate N's through a dipole portal can be divided into two categories.

(i) Inverse decays[2], $\nu + \gamma \to N$. These processes are important at $T \sim m_N$, and can be derived from the width of N.

(ii) $2 \to 2$ processes, such as $f^+ f^- \to N\bar{\nu}$ or $\bar{N}\nu$, where f is a SM fermion, as well as all crossing-related processes. While higher order in the coupling constant, these rates are enhanced in the UV.

At any given temperature in the early Universe, the abundance of N particles is set either by equilibrium, if their interaction rates are faster than the Hubble rate, or by the approach to equilibrium regulated by

$$\frac{n_N}{n_f} \sim \frac{\langle \sigma v \rangle n_{\bar{f}}}{H(T)}, \tag{18}$$

where n_f are n_N is the number density of charged species and HNLs, $H(T)$ is the Hubble rate, and $\langle \sigma v \rangle n_{\bar{f}}$ is the temperature-dependent rate for creating an HNL per unit of time. The most important for us is the scaling of the above expression with temperature and parameters of our model. Making a simple parametric estimate we arrive to

$$\frac{\langle \sigma v \rangle n_{\bar{f}}}{H(T)} \propto \alpha g_*^{-1/2} M_{\rm Pl} d^2 T, \tag{19}$$

where $M_{\rm Pl}$ is the Planck mass and g_* is the effective number of degrees of freedom appearing from the definition of the Hubble rate, $H(T) \simeq 1.66 g_*^{1/2} T^2 M_{\rm Pl}^{-1}$. The most important feature of Eq. (19), besides the self-explanatory dependence on $M_{\rm Pl}$ and d, is its scaling with temperature. The rate is enhanced in the UV, and therefore, it is *the highest temperatures* in the system that determine the initial abundance of N. Therefore, strictly speaking, one cannot determine the initial abundance of N without ever specifying the initial temperature relative to d^{-1}. On the other hand, assuming that the Universe at some point had temperature $T \sim d^{-1}$, the ratio in Eq. (19) is then larger than one for all values of d covered by our master plot, Fig. 1. Therefore, with this assumption, one can be sure that N was in fact thermalized in the early Universe.

Once N is thermally populated, it will last until the lifetime of the Universe is comparable to τ_N. To predict how much energy the thermally-created reservoir of N stores, one would need to understand at what temperatures HNLs decouple, which can be estimated parametrically by equating the r.h.s. of Eq. (19) to one. This gives the decoupling temperature of

$$T_{\rm decouple} \sim 1\,{\rm GeV} \times \frac{\tau_N}{0.1\,{\rm s}} \times \left(\frac{m_N}{10\,{\rm MeV}}\right)^3, \tag{20}$$

where we re-expressed d in terms of the lifetime formula for N. The decoupling of N means that at temperatures $T < T_{\rm decouple}$ the decays of heavy SM particles heat up the SM bath but not N, and its relative energy density is somewhat diluted as g_* at decoupling will be larger than at the time of decay. At the same time, for N heavier than an MeV, there is a possibility for a significant enhancement of the N energy density at decay due to them becoming nonrelativistic. The ratio $\rho_N/\rho_{\rm SM}$ will gain an enhancement factor $m_N/T_{\rm dec}$, where $T_{\rm dec}$ is the temperature corresponding to the time of the decay of N, $H(T_{\rm dec}) \sim \tau_N^{-1}$ (in the assumption that $T_{\rm dec} < m_N$). Consequently, our estimate becomes

$$\frac{\rho_N}{\rho_{\rm SM}} \propto \frac{g_N}{g_*(T_{\rm decouple})} \times \frac{m_N}{T_{\rm dec}}, \tag{21}$$

[2] The importance of inverse decays in astrophysical constraints of the neutrino dipole portals, including BBN and supernova bounds, was first discussed in [11].

where $g_N = 7/8 \times 4$ as N carries four fermionic degrees of freedom. This estimate can be used to constrain the lifetime of HNLs as $\rho_N/\rho_{\rm SM}$ is constrained at $T \sim 1\,{\rm MeV}$ through the n/p freeze-out. If $\tau_N \sim 0.1\,{\rm s}$, the ratio in Eq. (21) is $O(1)$, while only less than 10% variations are allowed (see, *e.g.* Ref. [60]).

B. Supernova SN 1987A

The modification of energy generation and transfer in stars can also serve to limit the viable parameter space for a dipole neutrino portal. In particular, SN 1987A has proved to be a useful probe of weakly coupled particles below the GeV scale [61–67]. The typical consideration is as follows: weakly coupled particles may serve to substantially enhance the rate of cooling of a supernova, and if this cooling proceeds too quickly and the energy is able to escape without being reabsorbed, then nuclear processes at the core of the supernova can rapidly stop. This in turn leads to significant deviations between the predicted and observed neutrino pulses observed at terrestrial neutrino observatories [68–70]. Therefore it is the rate of cooling, rather than the rate of production itself that is important.

There are two considerations in determining whether HNLs (or any new weakly coupled particle) can spoil supernova predictions. First, for sufficiently weak coupling very few HNLs will be produced, and consequently they will not be able to efficiently cool the interior of the supernova. This naively suggests strong couplings can be excluded, however, if the coupling is sufficiently large, then any HNLs that are produced will be trapped. Provided this trapping occurs within the "neutrinosphere" (defined as $r < R_\nu$ where $T(R_\nu) = 3$ MeV) [66], then the energy stored in the HNLs can be efficiently recycled and re-emitted in the form of neutrinos, ultimately having no impact on the observations at terrestrial detectors. A full treatment that captures this competition between production and absorption would involve a detailed study[3] of the following integrals [66]

$$\frac{{\rm d}E}{{\rm d}t} = \int_0^{R_\nu} P_{\rm esc}(r_0) \times \left\langle E_N \frac{{\rm d}\Gamma_{\rm prod}}{{\rm d}r} \right\rangle (r_0)\ {\rm d}r_0 \qquad (22a)$$

$$P_{\rm esc}(r_0) = \exp\left[- \int_{r_0}^{R_{\rm far}} \frac{1}{\lambda_{\rm MFP}} {\rm d}r \right]. \qquad (22b)$$

where ${\rm d}\Gamma/{\rm d}r$ is the local rate of production of HNLs, E_N denotes the HNL energy, $R_{\rm far}$ is a large radius to which the escape probability is insensitive, and the average is taken with respect to the local thermal bath at r_0. The probability of escape $P_{\rm esc}$ is found by exponentiating the line-of-sight integral of the mean free path, which in the case of the dipole portal will always be inversely proportional to the square of the dipole coupling $\lambda_{\rm MFP} \propto 1/d^2$.

For each HNL mass m_N, there will exist a minimal dipole coupling $d_{\rm prod}(m_N)$ for which too few HNLs are produced to significantly alter the observed neutrino signal. Likewise, there will also exist a maximum dipole coupling $d_{\rm abs}(m_N)$ such that for any stronger couplings the HNLs will be efficiently reabsorbed and will not cool the interior of the supernova appreciably. The region of excluded parameter space lies between theses two curves in the $d - m_N$ plane *i.e.* $d_{\rm prod}(m_N) < d_{\rm excl}(m_N) < d_{\rm abs}(m_N)$. Although Eqs. (22a) and (22b) are in general complicated, in the weak coupling regime ($d \lesssim d_{\rm prod}$), and the strong coupling regime ($d \gtrsim d_{\rm abs}$), the analysis simplifies.

In trying to obtain the lower curve $d_{\rm prod}(m_N)$ of Fig. 11, the coupling is small and so the probability of escape is nearly unity. We may therefore study the production of HNLs and neglect the absorptive properties of the bath. Furthermore, this may be done locally, as opposed to globally, at a characteristic radius. This approximation is often termed the "Raffelt criterion" [61], and is defined in terms of the energy carried by HNLs per unit volume, per unit time, ${\rm d}\mathcal{E}_N/{\rm d}t$ (being referred to as emissivity throughout this paper), at a fixed radius r_0

$$\frac{{\rm d}\mathcal{E}_N}{{\rm d}t} \leq 10\% \frac{{\rm d}\mathcal{E}_\nu}{{\rm d}t} \approx \left[\frac{\rho(r_0)}{{\rm g/cm}^3} \right] \times 10^{19} {\rm erg\ cm^{-3} s^{-1}}, \quad (23)$$

where ${\rm d}\mathcal{E}_\nu/{\rm d}t$ is the maximum energy per volume per time emitted via neutrinos. This criterion essentially requires that HNLs produced at some fixed radius r_0 carry no more than 10% of the total energy lost to neutrinos per time. The emissivity constraints derived based on the Raffelt criterion and from the criterion with the integrated energy are compared explicitly in [64]. The difference is well within an order of magnitude as demonstrated for their scenario.

In the limit of strong coupling, the relevant question is whether the produced HNLs can escape the supernova's neutrinosphere. Since $d_{\rm abs}(m_N) \gg d_{\rm prod}(m_N)$ we may assume a large flux of HNLs in the parameter space of interest, and so by Eq. (22a), it is the probability of escape that must inhibit cooling due to HNL production. As demonstrated by Eq. (22b), and the discussion thereafter, this quantity depends exponentially on the dipole coupling by way of the mean free path. Therefore, a reasonable criterion is that that $P_{\rm esc}(d_{\rm abs}) = 1/2$, since for $d \gtrsim d_{\rm abs}$ this quantity will be exponentially suppressed. Although the Raffelt criterion is most naturally imposed where the temperature is maximal, and densities are high, it is possible that this will lead to a rather conservative bound on $d_{\rm abs}$. This is because, being produced in the hot and dense interior of the su-

[3] Equation (22b) assumes an outward radial path for the HNL and does not account for passage through the core of the supernova. Neglecting this $\mathcal{O}(1)$ effect is already an approximation [66].

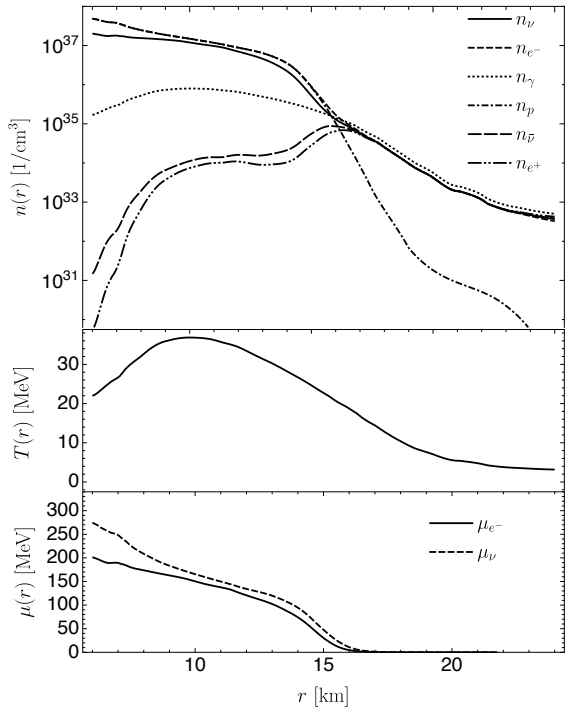

FIG. 10. Radial profiles of the number density, temperature, and chemical potentials at one second after the bounce from the simulation of an $18M_\odot$ progenitor [65].

pernova, the HNLs must travel through several kilometres of absorptive material composed of electrons, protons, and neutrinos, all of which have number densities in excess of $10^{37}/\text{ cm}^3$. This feature is mitigated to some extent due to Pauli-blocking, however which effect is dominant is hard to determine. With this in mind, we perform our analysis at two radii $r_0^{(a)} = 10$ km and $r_0^{(b)} = 14$ km. The former corresponds to the conventional choice [61, 62, 64, 66] of the hottest ($T \approx 30$ MeV) and most dense ($n_e, n_\nu, n_p \approx 10^{37}/\text{ cm}^3$) region of the supernova. The latter choice, by contrast, includes slightly lower temperatures ($T \approx 20$ MeV) number densities ($n_e, n_\nu, n_p \approx 10^{36}/\text{ cm}^3$) but does not require transit through the most dense regions of the supernova due to the sharp decline in number density in the outward radial direction.

Before turning to the details of the calculation of the emission rates and escape probabilities, we first summarize the physics that is included in our calculations. We use radial profiles corresponding to a supernova with an $18M_\odot$ progenitor, which are obtained by digitizing the reference runs shown in Fig. 5 of [65]. In calculating the optical depth, the full radial dependence is accounted for, but as discussed above, we apply the Raffelt criterion at two fixed radii. We include all species present except for neutrons as they do not couple to HNLs via the dipole portal. In computing the optical depth, and emissivities, we account for the effects of quantum degeneracy including Pauli-blocking, which is found to modify the rate of production and to have a dramatic effect on the escape probabilities of HNLs.

1. *Production*

Supernovae typically have significant populations of protons, neutrons and photons, as well as electrons and neutrinos, and their associated anti-particles. Save the neutron, HNLs couple to all of these species at tree level via the dipole portal, and this allows for the following production mechanisms

$$\nu + e^\pm \to N + e^\pm \qquad \text{(upscattering)} \qquad (24)$$
$$\nu + p \to N + p \qquad \text{(upscattering)} \qquad (25)$$
$$e^+ + e^- \to \overline{\nu} + N \qquad \text{(synthesis)} \qquad (26)$$
$$\gamma + \nu \to N \qquad \text{(inverse decay).} \qquad (27)$$

We point out that our analysis does not include thermal field theory effects, and so we omit the "plasmon decay" $\gamma \to \overline{\nu} N$ production mode. In general, ignoring the thermally acquired effective mass of photons in T channel scattering processes is only justified if the characteristic momentum flowing through the photon is much larger than its effective mass, which is on the order of $20 - 30$ MeV. Using vacuum propagators for the dominant HNL production process $e^-\nu \to e^- N$, we calculated the quantity $\sqrt{-\langle q^2 \rangle}$ and found it to be greater than 70 MeV for all masses considered, eventually asymptoting to m_N for heavy N. Furthermore, for all masses considered, ignoring the regime $\sqrt{-q^2} < 30$ MeV changes $\langle q^2 \rangle$ by less than 4%. In addition to thermal effects, we also neglect the influence of nucleon magnetic moments (because of the additional $\propto m_p^{-1}$ suppression), and for that reason neglect $\nu + n \to N + n$ production mode. Going back to the channels we consider, all of these have two incident species, and so the rate of production is controlled by the product of their densities (*i.e.* $n_e n_\nu$ in the case of electron upscattering). In the case of upscattering, however, the chemical potential can be an order of magnitude larger than the temperature, and so Pauli-blocking of the outgoing SM product must also be taken into account.

As discussed above, in considering the minimal dipole coupling that can spoil predictions from SN 1987A, we study the Raffelt criterion, Eq. (23), at both $r_0 = 10$ km and $r_0 = 14$ km.

The following integral equation defines the emissivity

$$\frac{\mathrm{d}\mathcal{E}_N}{\mathrm{d}t} = \int \frac{\mathrm{d}^3 p_1}{(2\pi)^3} \frac{\mathrm{d}^3 p_2}{(2\pi)^3} f_1 f_2 \langle E_N \sigma \rangle_{\mathfrak{F}} \, v_{\text{M}\o\text{l}}, \qquad (28)$$

where $f_a = 1/(\exp[(E_a - \mu_a)/T(r_0)] + 1)$ is the Fermi-Dirac distribution for species a, and $v_{\text{M}\o\text{l}}$ is the Møller

velocity

$$v_{\text{Møl}} = \sqrt{(\boldsymbol{v}_1 - \boldsymbol{v}_2)^2 - (\boldsymbol{v}_1 \times \boldsymbol{v}_2)^2}. \tag{29}$$

The average, $\langle E_N \sigma \rangle_{\mathfrak{F}}$, is taken over phase space with the appropriate distribution functions included. For inverse decays, this is the trivial one-body phase space of the HNL, but for $2 \to 2$ process the appropriate Pauli-blocking factor of the outgoing SM particle, $\mathfrak{F}(E_3) := 1 - f(E_3)$, is included, where E_3 is evaluated in the rest frame of the bath. Explicitly, for $2 \to 2$ processes the average is defined as

$$\langle E_N \sigma \rangle_{\mathfrak{F}} := \int \frac{\mathrm{d}\Phi_2(p_3, p_N)}{4\mathcal{F}(s)} \mathfrak{F}(E_3) E_N |\mathcal{M}|^2_{\text{prod}} \tag{30}$$

where $\Phi_2(p_3, p_N)$ denotes the two-body Lorentz invariant phase space of the outgoing HNL and SM particles, $\mathcal{F}(s)$ the Lorentz-invariant flux factor, and E_N, like E_3, is evaluated in the rest-frame of the bath. The production matrix element $\mathcal{M}_{\text{prod}}$ is calculated at zero-temperature, and does not include—for example—the in-medium modification of the photon propagator.

Following [62, 71, 72], we can rewrite Eq. (28) as

$$\begin{aligned} \frac{\mathrm{d}\mathcal{E}_N}{\mathrm{d}t} =& \frac{1}{32\pi^4} \int_{M^2}^{\infty} ds \int_{\sqrt{s}}^{\infty} dE_+ \int dE_- \langle E_N \sigma \rangle_{\mathfrak{F}} \\ & \times F(s, m_1, m_2)\ f(E_1, \mu_1)\, f(E_2, \mu_2) \end{aligned} \tag{31}$$

where

$$\begin{aligned} M^2 &= \text{Max}[(m_1 + m_2)^2, (m_N + m_{\text{final}})^2], \\ E_1 &= \frac{E_+ + E_-}{2} \quad \text{and} \quad E_2 = \frac{E_+ - E_-}{2}, \\ F(s, m_1, m_2) &= \sqrt{\frac{1}{4}\left(s - m_1^2 - m_2^2\right)^2 - m_1^2 m_2^2}. \end{aligned} \tag{32}$$

Using the Mandelstam variable s, we can show that E_- depends on s, E_+, m_1, m_2, and $\cos\theta$, and that its associated bounds of integration are obtained by considering the limits $\cos\theta \to \pm 1$ with E_+ and s held fixed.

2. *Escape*

The escape probability Eq. (22b) is dictated by the mean free path λ_{MFP} of the HNL in the hot bath of the supernova. Demanding that the probability of escape is less than 50% is equivalent to demanding that $-\ln P_{\text{esc}} \lesssim 2/3$. Since the dipole portal is the only coupling between the Standard Model and the HNL, all processes that contribute to λ_{MFP} are proportional to d^2. It is therefore convenient to introduce a reduced mean free path $\lambda\!\!\!^{-}$, defined at a reference value $d = 10^{-7}$ GeV^{-1} via

$$\begin{aligned} -\ln P_{\text{esc}} &= \int_{r_0}^{25\text{ km}} \frac{1}{\lambda_{\text{MFP}}(r)} \mathrm{d}r \\ &= \left(\frac{d}{10^{-7}\text{ GeV}^{-1}}\right)^2 \int_{r_0}^{25\text{ km}} \frac{1}{\lambda\!\!\!^{-}_{\text{MFP}}(r)} \mathrm{d}r \end{aligned} \tag{33}$$

Implicit in the above analysis is the assumption that the path of the HNL is directed radially outwards. This underestimates the probability of absorption as it neglects paths that travel through the core and other overdense regions, however as discussed in Appendix B. of [66] this effect is $\mathcal{O}(1)$ and can be captured by multiplying the optical depth by the substitution $\lambda\!\!\!^{-}_{\text{MFP}} \to \lambda\!\!\!^{-}_{\text{MFP}}/3$. We may then define the critical dipole moment where HNLs are efficiently trapped via the condition

$$d_{\text{abs}} = \sqrt{\frac{2/3}{3 \times \int_{r_0}^{25\text{ km}} \frac{\mathrm{d}r}{\lambda\!\!\!^{-}_{\text{MFP}}(r)}}} \times 10^{-7}\text{ GeV}^{-1}. \tag{34}$$

The above procedure does not take into account the flux of HNLs coming from the core of the supernova can be exponentially large, and therefore some of amount of energy deposition can happen beyond d_{abs}. The flux is a factor of $(d_{\text{abs}}/d_{\text{prod}})^2 \sim 10^6$ larger than the lower bound, and so an even larger dipole coupling is required to efficiently absorb this large flux of HNLs, given roughly by $d_{\text{abs}} \to d_{\text{abs}} \times \log\left(\frac{d_{\text{abs}}}{d_{\text{prod}}}\right)^2$, which is approximately an order of magnitude larger, and consequently more stringent. Since we neglect this effect, our analysis can be considered conservative in this regard.

Both single body decay of the HNL, and $2 \to 2$ scattering contribute to the mean free path. In no particular order, the relevant processes are

$$N + e^{\pm} \to \nu + e^{\pm} \qquad \text{(downscattering)} \tag{35}$$

$$N + p \to \nu + p \qquad \text{(downscattering)} \tag{36}$$

$$\overline{\nu} + N \to e^+ + e^- \qquad \text{(annihilation)} \tag{37}$$

$$N \to \gamma + \nu \qquad \text{(decay)} \tag{38}$$

$$N + \text{SN} \to N + \text{SN} \qquad \text{(gravitational trapping).} \tag{39}$$

We have included the full radial dependence of the temperature and chemical potentials in our calculation of Eq. (34). As can be clearly seen in Fig. 10, the chemical potentials of the neutrinos and electrons are significantly higher than the temperature within the interior of the supernova, therefore for HNLs produced at $r_0 \approx 10$ km, Pauli-blocking and the Fermi-Dirac distributions of the absorptive species can play an important role in determining the escape probability. As discussed above, we compute the reduced optical depth integral at a reference dipole coupling of 10^{-7} GeV^{-1} and include the effects of Pauli-blocking via

$$\frac{1}{\lambda\!\!\!^{-}_{\text{MFP}}(r)} = \sum_{\alpha} \langle n_\alpha \sigma_{\alpha N} \rangle(r) + \beta \gamma(r_0) \langle \Gamma_N \rangle(r). \tag{40}$$

Here $\alpha \in \{e^-, e^+, \nu_{e,\mu,\tau}, \overline{\nu}_{e,\mu,\tau}, \gamma, p\}$ labels the species that can absorb HNLs, and n_α's are their Fermi-Dirac distributions. The thermal averages $\langle n_\alpha \sigma_{\alpha N} \rangle$ and $\langle \Gamma \rangle$ includes the thermal distribution of the absorptive bath for the $2 \to 2$ absorption, and the associated Pauli-blocking of outgoing SM particles for both decays and $2 \to 2$ processes.

We fix the incident HNL energy to be $\langle E_N\rangle(r_0, m_N)$, defined as the average energy per HNL produced at $r_0 = 10$ or 14 km, and this implies a boost factor for the HNL $\beta\gamma(r_0, m_N)$. In practice we compute the average energy numerically, however the qualitative behaviour can be understood as follows. The dominant production mechanism over most of the mass-range is Primakoff upscattering off of electrons which is Pauli-blocked on the outgoing electron. For $m_N \gtrsim \mu_e$, the momentum transfer required to create an HNL typically kicks electrons above the Fermi surface and imparts the HNL with three-momentum of order $P_N \sim \mathcal{O}(\mu_e)$. Therefore the momentum can be estimated using elementary kinematics. In contrast, for low masses the effects of Pauli-blocking must be accounted for by demanding a large momentum transfer, $q^2 \approx -\mu_e^2$, so as to kick the electron above the Fermi-surface. Taking the average neutrino to be $\mu_\nu/2$ and averaging over angles then leads to the estimate

$$\langle E_N\rangle \approx \begin{cases} m_N + \frac{\mu_e^2}{2m_N} & \text{for } m_N \gg \mu_e \\ \frac{m_N^2}{\mu_\nu} + \mu_e & \text{for } m_N \ll \mu_e \end{cases}, \qquad (41)$$

where the chemical potentials are evaluated at r_0. We also assume the HNL's path is directed radially outward (and correct for the possibility of transit through overdense regions via a factor of 3 as discussed above).

The thermal averages $\langle n_\alpha \sigma_{\alpha N}\rangle$ and $\langle \Gamma_N\rangle$ take into account the radial profile of the supernova, as a consequence of the Pauli-blocking of outgoing SM particles and the thermal distributions of initial SM particles inheriting the radial dependence of the chemical potential and temperature profiles. As in the case of production, the matrix element $|\mathcal{M}_{\rm abs}|^2$ is computed at zero temperature and we have checked that finite temperature corrections are under control.

Finally, the gravitational pull from the supernova could potentially trap the HNLs and prevent additional cooling of the supernova from happening. This is especially relevant for the high mass regime. Here we follow the simple energy argument introduced in [73] that determines the particle mass for which this effect becomes important.

The gravitational trapping has to be taken into account when

$$\langle E_{\rm kin}\rangle_{\rm HNL} \leq \frac{G M_c m_N}{R_c}, \qquad (42)$$

where $\langle E_{\rm kin}\rangle_{\rm HNL}$ is the average kinematic energy of the HNLs, G is the Newton constant, M_c is the enclosed mass of the supernova within the radius R_c, at which the HNL of mass m_N is produced. We take $M_c \approx M_{\rm SN}$ which is the mass of SN 1987A and calculate $\langle E_{\rm kin}\rangle_{\rm HNL}$ at two radii R_c =10 km and 14 km, corresponding to the radii we choose for the emissivity and the optical depth considerations. We determine that for $m_N \gtrsim 320$ MeV gravitational trapping is important at both $R_c = 10$ km and 14 km.

C. Results

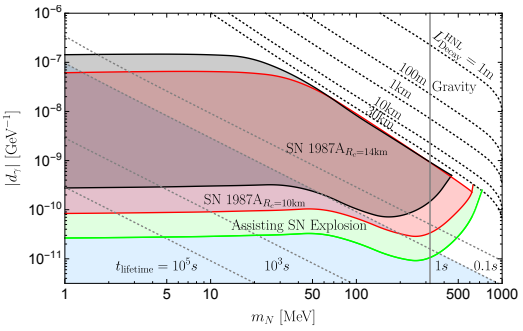

FIG. 11. Emissivity and optical depth constraints (red) from supernovae SN 1987A, and parameter space facilitating its conversion to a neutron star (green). We also show lines of constant HNL lifetimes to gauge where BBN might be affected. Two radii of production r_0 are plotted for comparison, with one at the hottest densest radius $r_0 = 10$ km and one closer to the edge of the high density region $r_0 = 14$ km. The gravitational trapping becomes significant for HNLs with mass above the vertical gray line, labeled "Gravity".

We begin with the BBN limits, that rest on several assumptions. First, we assume that the temperatures in the early Universe were initially rather large, and as a consequence, HNLs got thermally populated. If the maximum (*i.e.* reheating after inflation) temperature was limited to a sub-GeV range (which is a rather extreme assumption), then domains of parameter space with small m_N and small d will not be constrained by n/p freeze-out, as the abundance of HNLs at 1 MeV can be much smaller. The second assumption is that we assume that the BBN proceeds along a standard scenario, and HNLs provide only a small perturbation. An alternative scenario, when the Universe is actually dominated by N, and its decay reheats the ν and γ, e baths, might not be excluded throughout the whole parameter space. Namely, the BBN provides only a handful of reliable predictions (^{4}He, D/H). It could be possible that for some "islands" on $\{m_N, d\}$ space, the outcome of the nuclear reaction network is similar to a standard BBN. In this case, however, one would also have to make sure that the energy densities of neutrinos and photons are also consistent with measurements of $N_{\rm eff}$. This may look as an additional fine tuning, and therefore we do not consider such an accidental possibility seriously.

Thus, with the above caveats, if the $\rho_N/\rho_{\rm SM}$ ratio is larger than 0.1 at the time of n/p freeze-out, the BBN is perturbed outside of its agreement with observations. Then it is possible to set the constraints on lifetime to be less than a fraction of a second (see Ref. [74] for a somewhat similar analysis of the Higgs portal relics). We choose to be on a very conservative side, and set the limit for lifetime to be 1 sec, shown by the diagonal line in Fig. 11. (At $m_N \sim 1$ MeV, the decoupling temperature

is close to an MeV, and therefore $\rho_N/\rho_{\rm SM} > 0.2$ unless N particles decay early. At $m_N > 10$ MeV, the decoupling temperatures are in the GeV regime and larger, so that there can be a significant dilution by $g_*(T_{\rm decouple})$. However, $m_N/T_{\rm dec}$ more than compensate for this dilution, along the $\tau_N = 1$ second line). We observe that on $\{m_N, d\}$ space the BBN constraints do not overlap with neutrino/beam dump or high energy experimental constraints.

Our astrophysical results are collected in Fig. 11. As described in detail in the previous subsections, we have calculated present limits on heavy neutral lepton dipole moments stemming from supernovae cooling. The lower curve of the excluded region is found by requiring that the rate of energy produced by HNL (the emissivity) is larger than a tenth of that from neutrinos. The upper curve is obtained by enforcing that $\int \lambda_{MFP}^{-1} {\rm d}r < 2/3$, namely that the probability of an HNL interacting with something on its way outside the star (the optical depth) is small.

Our analysis reveals that Pauli-blocking of electrons and neutrinos is an essential feature in determining both the emissivity and especially the optical depth. In the latter case, quantum degeneracy makes the hot and dense interior of the supernova nearly transparent to HNLs whose decay and downscattering is inhibited by a Fermi-sea extending up to momenta on the order of $\mu_\nu \approx 250$ MeV. Unintuitively, this means that the escape probability for an HNL produced at $r_0 = 10$ km is nearly equal to that of one produced at the edge of the densest regions at $r_0 = 14$ km. Similarly, the production of HNLs is severely inhibited by the Fermi-sea of electrons. Naively, the high densities of electrons and neutrinos shown in Fig. 10 favor HNL production, and this suggests that Primakoff upscattering is the dominant production mechanism. This is, in fact, the case at low masses (but only marginally so), however at higher masses ($m_N \gtrsim 50$ MeV) inverse decays actually come to dominate despite the number density of photons being two orders of magnitude smaller. This is because the inverse decay is not Pauli-blocked. The consequences of quantum degeneracy are that the HNL behaves as if it is much more weakly coupled than one would expect based on naive predictions.

The qualitative features of our results can be described as follows. The upper curve is dominated at low masses by downscattering off of electrons and neutrinos, and the inclusion of Pauli blocking increases the bound on d due to the large chemical potentials (i.e. a large number of already occupied states) of these leptons for $r \geq 10$ km. Downscattering is relatively insensitive to the mass of the HNL, (*i.e.* $\sigma \sim d^2$) and so is eventually overtaken by the decay of the HNL which scales as $\Gamma \propto d^2 m_N^3$ and benefits from the absence of Pauli-blocking on the outgoing photon; this crossover between m_N independent downscattering, and power-law decay lengths can be clearly seen in Fig. 11. The bottom curve is dominated primarily by upscattering of neutrinos off of electrons. This process is only Pauli-blocked on the outgoing electron, and benefits from high number densities of both electrons and neutrinos. In direct parallel with the escape probabilities, this process is eventually overtaken at large masses by inverse decays. The inverse decays scale as $m_N^4 d^2$ and provide the dominant contribution for $m_N \gtrsim 50$ MeV. The maximal emission is reached when $m_N \simeq \sqrt{s} \approx (T + \mu_\nu)$, but this production channel ceases to be viable at masses much higher than the average center of mass energy $m_N \gg \langle\sqrt{s}\rangle \approx (T + \mu_\nu) \approx 250$ MeV because the HNL cannot be efficiently produced. Upscattering has a slightly higher kinematic limit of $m_N \gg \langle\sqrt{s}\rangle \approx (\mu_e + \mu_\nu) \approx 500$ MeV due to the large chemical potential of the neutrinos.

Gravitational trapping of the HNLs becomes important for large mass HNLs. Above the mass $m_N = 320$ MeV, the average kinematic energies of the HNLs are smaller than gravitational potential they feel from SN 1987A, as indicated with a vertical line in Fig. 11. The effect can to some degree alleviate the cooling bound of the SN on the HNLs since these HNLs can be gravitationally trapped and never travel out of the supernova. We leave a more refined determination of the gravitational effect on the SN cooling to future works.

Also on the plot there is a region called "Assisting SN Explosions". The detailed mechanism of core-collapse supernova explosion is an active research topic, and with the most explored mechanism being driven by neutrinos [75]. Simulation results such as [76, 77] have tended to find that the neutrino-driven explosion struggles to reproduce the revival of the shock-waves for a successful explosion, and requires additional shock energy to match the observation during the core collapse. It is worth noting, however, that the most recent simulation based on a 3D progenitor model [78], suggests that the neutrino-driven mechanism itself could possibly provide enough shock revival and explain the observed explosion energies. It is likely that a larger range of progenitors and more refined simulations are still required to fully understand the issue of SN explosions.

With these details in mind, it is worth noting that new degrees of freedom, for example, HNLs, have long been proposed to power SN explosions [79], and were most recently proposed to assist neutrinos in reviving the shock waves and augment their energies [80]. We briefly review the mechanism for the reader. The star begins by collapsing under its gravitational pull, causing a bounce off of the inner core. This radiates an outward shock. The shock gets stuck, because of dissociation of heavy nuclei, and gets revived by SM neutrino heating and hydrodynamic effects, producing an explosion. This depletes the star's core of leptons. The outward shock then encounters a matter envelope surrounding the star. At this point, previous simulations [76, 77] found that the shockwave is not able to expel the envelope, and the explosion is quenched. The matter in the envelope falls back into the core, possibly creating a black hole and preventing a neutron star final state from forming. If it was blown away,

however, the core could live on as a neutron star, which is the observed remnant of the core-collapse supernovae. By adding HNLs (or any other metastable particles with right properties), they can escape to the envelope and decay into neutrinos and photons. This creates an additional outward radial pressure in the envelope and breaks up some of the heavy nuclei. The original shockwave then has an easier time expelling the envelope away, and wastes less energy dissociating the nuclei inside the envelope. Interestingly, even a small amount of additional energy injection could possibly result in a proper explosion [79, 80].

In Fig. 11, the upper bound of the 'preferred' region for assisting supernova explosions is determined by consistency with SN 1987A limits. The lower bound, which is the main numerical result of [80], corresponds to having an energy emission from HNLs of 10^{51}ergs. By contrast, the energy emitted by all Standard Model neutrino species in SN 1987A is $E_\nu \approx 3 \times 10^{53}$erg [61]. It is important to note that the simulation in [80] assumes vacuum flavor mixing angles of $\sin^2\theta > 10^{-8}$ for ν_τ mixings and $10^{-8} < \sin^2\theta < 10^{-7}$ for ν_μ mixings, which are not present in our model. However, the main features of their analysis still hold in our case, since the HNLs in our scenario can also generate the required amount of energy injection given in [80]. To obtain the favored region, we have effectively redone the emissivity analysis described in Section V B 1 using an emitted energy of 10^{51}erg. Recall that the emissivity constraint is done requiring a power loss through HNLs less than 10% that from SM neutrinos. We find that the "Assisting SN Explosion" regime is mostly covered by the BBN constraint on HNLs with lifetimes longer than 1 second.

VI. DISCUSSION AND CONCLUSION

In this paper we have considered a variety of phenomenological consequences of a massive Dirac particle, that has a dipole portal d to the SM neutrinos and the photon, as a main source of production and decay of HNLs. The Dirac nature of the mass of N is dictated by the arguments of the neutrino mass generation. Different variants of such models have been proposed in the past, as a way of mimicking the excess of neutrino signals observed at LSND and MiniBooNE. We have provided an attempt at a comprehensive analysis of this class of model, assuming the dominance of dipole couplings.

We find that the high energy probes (LEP and LHC) of HNLs through a dipole portal are giving sensitivity to d at a scale of $(10\,\text{TeV})^{-1}$ and better, mostly through the mono-photon type signatures. In particular, the sensitivity of the LHC experiments extends to the TeV scale m_N. High intensity beam dump and neutrino experiments ("intensity frontier" experiments) cannot reach to such high masses, but instead are able to probe much lower values of couplings for the sub-GeV masses. We find that the inclusion of the dipole production of N disfavors common explanation of the MiniBooNE and LSND anomalies by already existing data. Interestingly, LSND itself provides the most stringent constraints on the dipole coupling at low masses, while the MiniBooNE, MicroBooNE, and SBND detectors provide the leading constraints at slightly higher masses. At the peak sensitivity to the dipole coupling, for $m_N \sim$ few 100 MeV, the experiments probe scales of $d \sim (10^{-7} - 10^{-6})\,\text{GeV}^{-1}$, which is far beyond the weak scale. Future experimental facilities, including SBND, and in particular SHiP, will be able to help improve sensitivity to these couplings. For the SHiP main detector, the level of the single photon backgrounds is not currently well understood, and while we use our optimal estimates at this point, detailed simulations can help better evaluation of sensitivity to dipole portal. Astrophysics, in particular physics of SN explosions, further restricts the parameter space for the model, probing up to a few hundred MeV scale masses and a $d \sim (10^{-7} - 10^{-10})\,\text{GeV}^{-1}$ range of couplings. The cosmological bounds are somewhat model dependent as they are sensitive to the high-temperature regime of the early Universe for which we do not have the direct experimental data. In the most likely eventuality of high initial temperatures, the constraints on lifetime are in the 1 second range and better, disfavoring low-d, low-m_N corner of the parameter space. Overall, the HNL coupled to the dipole portal adds to new physics models that can be studied both at high and medium energies, and in astrophysical/cosmological settings. We conclude our paper with a few additional comments:

(i) One of the reasons the current model can be studied with such a variety of tools is the fact the dipole portal we explore, below the electroweak scale, is a dimension 5 portal. It gives cross sections that scale as $\sigma \propto d^2$. This is similar to the interactions of axion-like particles a (*e.g.* $g_{a\gamma\gamma} a F\tilde{F}$), which is also dimension 5. Indeed, one can observe broad numerical similarities between sensitivity to $g_{a\gamma\gamma}$ and our derived sensitivity to d.

(ii) We have covered only a handful of the existing intensity frontier searches that we think to be the most sensitive. It is possible that some other experiments (such as *e.g.* CHARM, CCFR, and T2K) may also provide additional constraints on the model. Among new planned facilities, some would involve unprecedented intensities (DUNE), and it is possible that new levels of sensitivity to d can be derived there as well.

(iii) There are several experimental setups proposed at the LHC to probe long-lived particles, including MATHUSLA [81] and CODEX-b [82], and a small detector to probe weakly coupled states in the forward regime, FASER [83], which has already considered HNLs (but not their neutrino dipole interactions) [84, 85]. These setups could potentially extend our reach at the energy frontier. However,

since the lifetime of the HNLs in our scenario scales as m_N^{-3} as seen in Eq. (6), the decay lengths may be too short in the near GeV mass range to significantly improve on the reach of existing probes.

(iv) We have provided a SM gauge invariant completions of the dipole portal operator. This should not be confused with a proper UV completion, which was briefly discussed in [11, 12]. See also [26] for a one-loop calculation of HNL radiative decay rates in the context of various renormalizable electroweak gauge theories. Such UV completion may also point to a potential tuning issue that can arise in this model. Operators (10) can radiatively induce significant mass mixing operator, LNH, which we have assumed to be small and/or absent. It will be important to find out whether tuning-free UV completions of this model exist. This task falls outside the scope of this paper.

(v) The setup considered in this paper can be easily extended to provide new constraints and unprecedented sensitivity reaches for other weakly interacting particles, including millicharged particles [86].

ACKNOWLEDGEMENTS

We thank Ornella Palamara, Roni Harnik, and Zarko Pavlovic for the correspondence and for the details of the experiments along the Booster Neutrino Beamline at Fermilab. We also thank Jae Hyeok Chang and Robert Lasenby for useful references in astrophysical and cosmological constraints, and Jonah Miller for discussions relating to supernova simulations. In addition, we thank Cliff Burgess, Pilar Hernandez, Richard Hill, Joseph Bramante and Christopher Brust for useful discussions. This research was supported in part by Perimeter Institute for Theoretical Physics. Research at Perimeter Institute is supported by the Government of Canada through the Department of Innovation, Science and Economic Development and by the Province of Ontario through the Ministry of Research and Innovation. This research was also supported by funds from the National Science and Engineering Research Council of Canada (NSERC), and the Ontario Graduate Scholarship (OGS) program. YT was supported by the Visiting Graduate Fellow program at Perimeter Institute.

Appendix A: Intensity Frontier

1. Neutrino upscattering

We obtain an expression for $d\sigma/dt$. Consider first the matrix element for the production of N, which factorizes into a hadronic and a leptonic tensor, i.e.

$$|\mathcal{M}|^2 = \frac{|d|^2 e^2}{q^4} L_{\mu\nu} W^{\mu\nu}. \tag{A1}$$

In terms of a right-handed projection operator, the leptonic tensor is

$$L_{\mu\nu} = 4\,\mathrm{Tr}\Big[\not{p}_1 P_R\, \sigma_{\nu\alpha} q^\alpha (\not{p}_3 + m_N) \sigma_{\mu\beta} q^\beta\Big]. \tag{A2}$$

The hadronic current is given by

$$\langle A|\Gamma^\mu|A'\rangle = F_1\gamma^\mu + F_2 \frac{i}{2M_H}\sigma^{\mu\delta} q_\delta. \tag{A3}$$

In the heavy nucleus limit, squaring Eq. (A3) gives

$$W^{\mu\nu} = F_1^2 \mathrm{Tr}\Big[\left(\not{p}_4 + m_H\right)\gamma^\mu\left(\not{p}_2 + m_H\right)\gamma^\nu\Big]. \tag{A4}$$

The representation of the form factors will depend on whether the scattering is coherent or inelastic. In the former case, the neutrino upscatters on the nucleus as a whole and the cross section scales as Z^2. Since $M_H = AM_{\mathrm{nucleon}}$ and $|t| = |q^2| = Q^2$ is small, we retain only F_1 in Eq. (A3), which we take to be the Woods-Saxon (WS) form factor. Indeed, the contribution of the magnetic moment of a nucleus to the $\nu \to N$ transition is relatively small, especially for large A nuclei, for the case of coherent scattering. In the shell model description, only the outside shell nucleons contribute to nuclear spin and magnetic moment. Therefore, the magnetic moment provides only a very subdominant $\mathcal{O}(1/A)$ part of the amplitude relative to the main (Coulomb-induced) part. For small A nuclei, this may be more important. Note that when the scattering occurs at larger Q^2, the magnetic moments of the nucleons are being taken into account. The WS form factor parameterizes the charge density of the nucleus as

$$\rho(r) = \frac{\rho_0}{1 + \exp\left(\frac{r - r_0 A^{1/3}}{a}\right)} \tag{A5}$$

and takes its Fourier transformation with respect to the momentum exchange q [87, 88]. From Eq. (A1), we obtain

$$\begin{aligned}\frac{d\sigma}{dt} = -\frac{2\alpha|d|^2 Z^2 F_{\mathrm{WS}}^2}{t^2\left(s - m_H^2\right)^2}\times \\ \Big\{ -tm_N^2(2s+t) + m_N^4\left(2m_H^2 + t\right) + \\ 2t\left(s - m_H^2\right)\left(-m_H^2 + s + t\right)\Big\}\end{aligned} \tag{A6}$$

The $1/t^2$ pre-factor in the lab frame is proportional to $1/(E_N - E_\nu)^2$, meaning there is a phase space enhancement favoring $E_N = E_\nu$.

On the other hand, when the scattering is inelastic, the incoming neutrinos scatter off of the individual nucleons. When this happens, $|t|$ is of moderate size,

$M_H = M_{\rm nucleon}$ and we retain both form factors. F_1 and F_2 take on different values depending if they are for the neutron or proton. Their values are given [89, 90] by solving the system of equations

$$\begin{aligned} G^{p,n}_{\gamma,E} &= F^{p,n}_{1,\gamma} - \frac{Q^2}{4M^2_{\rm Nucleon}} F^{p,n}_{2,\gamma} \\ G^{p,n}_{\gamma,M} &= F^{p,n}_{1,\gamma} + F^{p,n}_{2,\gamma} \end{aligned} \tag{A7}$$

with

$$\begin{aligned} G^{\{p,n\}}_{\gamma,E} &= \{G_D, 0\} \\ G^{\{p,n\}}_{\gamma,M} &= \mu_{\{p,n\}} G_D \\ G_D &= \frac{1}{(1 + Q^2/0.71\ {\rm GeV}^2)^2} \\ \mu_{p,n} &= \{2.793, -1.913\}. \end{aligned} \tag{A8}$$

We then obtain

$$\sigma_{total} = Z \times \sigma_p + (A - Z) \times \sigma_n. \tag{A9}$$

In contrast to the coherent scattering case, the inelastic cross section depends only linearly on Z and A. Furthermore, values of t for which we have inelastic scattering generically avoid the $t \to 0$ enhancement.

2. Meson decays

In determining the number of HNLs present at intensity frontier experiments, it is important to consider both Primakoff upscattering and direct decays of mesons into HNLs. The decay in flight of mesons will lead to a distorted spectrum of HNLs that depends on the details of the decay at rest, and the spectrum of incoming mesons. In this appendix we outline how to obtain the spectrum of HNLs given a spectrum of incident mesons.

We denote the rest frame energy and momentum $\mathcal{E}$ and $\mathcal{P}$, and the angle relative to the boost vector in the rest frame as ϕ, while lab frame quantities are defined analogously as E, P, and θ. We first compute the rest frame differential decay rate as a function of the energy of the HNL ${\rm d}\Gamma/{\rm d}\mathcal{E}$. Normalizing by the overall decay rate of the meson defines the differential branching ratio in the rest frame ${\rm dBR}/{\rm d}\mathcal{E}|_{\rm rest} = (1/\Gamma) \cdot {\rm d}\Gamma/{\rm d}\mathcal{E}|_{\rm rest}$. The most important contribution to HNL production is from pions, and so we quote the result of ${\rm d}\Gamma/{\rm d}\mathcal{E}$ in the rest frame for the process $\pi^0 \to N\nu\gamma$

$$\frac{{\rm d}\Gamma}{{\rm d}\mathcal{E}} = -\frac{1}{2\pi m_\pi}\alpha^2 |d|^2 F_\pi^2 \left[\mathcal{P}\left(4\mathcal{E}^2 m_\pi^2 - 3\mathcal{E} m_\pi m_N^2 - 2\mathcal{E} m_\pi^3 + m_N^4 + 3 m_\pi^2 m_N^2\right) - m_\pi^2 \left(4\mathcal{E} - m_\pi\right) m_N^2 \tanh^{-1}\left(\frac{\mathcal{P}}{\mathcal{E}}\right)\right], \tag{A10}$$

In this expression, we use the notation $F_\pi = (4\pi f_\pi)^{-1}$ with $f_\pi \approx 92$ MeV. In our meson calculations, we have set the lepton masses in some of the integration bounds to 0 in order to make the integrals tractable. For most of the meson decay channels, this approximation was found to have a minor effect on the results. For heavy m_N, the $\pi \to \mu N\gamma$ channel was found to be underestimated by this approximation, yielding a conservative estimate. Next, for a given energy $\mathcal{E}$, the resultant distribution in the lab-frame can be found by considering

$$E = \gamma\mathcal{E} - \beta\gamma\mathcal{P}\cos\phi \tag{A11}$$

and noting that the decay of a pseudo-scalar is isotropic in the rest frame. Consequently the lab energies are sampled uniformly from $[E_-, E_+]$ where $E_\pm = \gamma\mathcal{E} \pm \gamma\beta\mathcal{P}$. The population of the interval of phase space in the lab frame must be the same as its corresponding interval in the rest frame. This implies that a delta-function distribution in the rest frame is transformed to a box distribution with a width of $(E_+ - E_-) = 2\gamma\beta\mathcal{P}$ in the lab frame.

The same argument can be applied to obtain the maximum and minimum rest frame energies that can be boosted into a given infinitesimal window centered about E. These are given by

$$\mathcal{E}_\pm = \gamma E \pm \beta\gamma P. \tag{A12}$$

Using this information we can construct the spectrum of HNL energies generated by a meson traveling at velocity β in the lab frame

$$\left[\frac{1}{\Gamma}\frac{{\rm d}\Gamma}{{\rm d}E}\right]_{\rm lab} = \int_{\mathcal{E}_A}^{\mathcal{E}_B} \frac{1}{2\gamma\beta P(\mathcal{E})}\left[\frac{1}{\Gamma}\frac{{\rm d}\Gamma}{{\rm d}\mathcal{E}}\right]_{\rm rest} {\rm d}\mathcal{E} \tag{A13}$$

where the factor of $2\gamma\beta\mathcal{P}$ accounts for the normalization of the box distribution discussed above. The quantities E_A and E_B are defined via $\mathcal{E}_A(E,\gamma) = \min(\mathcal{E}_-, \mathcal{E}_{\rm min})$ and $\mathcal{E}_B(E,\gamma) = \max(\mathcal{E}_+, \mathcal{E}_{\rm max})$ where $\mathcal{E}_{\rm min}$ and $\mathcal{E}_{\rm max}$ are the minimum and maximum energies of the HNL that are kinematically allowed in the rest frame. Notice that the limits of integration on the right-hand side are functions of the lab energy E and the velocity β, or equivalently γ.

Finally, we consider a spectrum of parent mesons. In this case a spectrum (e.g. $N(\gamma) = N(E/m_\pi)$ in the case of pions) is assumed to be given and we weight the contribution of each value of β by this spectrum finally giving

$$N_{\rm lab}(E_N) = \int_{\gamma_{\rm min}}^{\gamma_{\rm max}} \left[\frac{1}{\Gamma}\frac{{\rm d}\Gamma}{{\rm d}E}\right]_{\rm lab} N(\gamma){\rm d}\gamma \tag{A14}$$

the spectrum of HNL's produced from a given flux of mesons.

The meson energy lab spectrum used was adjusted to account for the magnitude of the beam energy, and the meson masses under consideration. When considering SBN the Sanford-Wang [32, 91] distribution was used to model the incident pions, while for kaons and eta mesons the Feynman scaling hypothesis [32] was employed. At SHiP where the incident proton beam has an energy of 450 GeV the BMPT [91, 92] distribution was used instead for both pions and eta mesons. The use of the Feynman scaling approach was inspired by Ref. [93], which argues that low energy proton beams and high meson masses exhibit special mass effects that are not well captured by Sanford-Wang. The Feynman scaling approach assumes that $\frac{d^2\sigma}{dpd\Omega}$ depends only on p_T and $x_F = p_{||}^{\rm COM}/p_{||}^{COM,max}$, and is proportional to $(1-|x_F|)$. Mass effects tend to give stronger weight in the data in the $x_F = 0$ regime. This is reflected in the Feynman Scaling approach, whereas Sanford-Wang keeps increasing as x_F crosses over to negative values. At even lower energies, such as at LSND where the POT energy is around 0.8 GeV, we employ the Burman-Smith distribution [37, 38]. By fitting to datasets spanning a wide range of pion kinetic energies ($30 - 553$ MeV) the Burman-Smith distribution attempts to model the pion spectrum down to zero kinetic energy. At LSND, as low kinetic energy protons interact with the beam stop, pions which are produced are slowed down. The negative pions are absorbed in matter while the positive pions decay. Most of these π^+ are at rest, while some (2%) decay in flight. For μ^+ and π^+ that decay at rest, we take their spectrum to be isotropic. For π^0 and π^+ that decay in flight, we use the Burman-Smith distribution.

3. Perturbative electroweak backgrounds

As another source of background, we consider non-resonance induced single photons from perturbative electroweak processes. Although it is intuitive that the loop suppressed SM backgrounds from $A\nu \to A\nu\gamma$ will be low, it is important to quantify by how much, as this process could occur via neutrinos interacting in the walls of the SHiP experiment. Our goal here is to show that this potential source of background is very small and under control. The cross section for $\gamma\nu \to \gamma\nu$ has been explicitly calculated using effective operators [94, 95], and this provides a convenient way to calculate the SM contribution to $A\nu \to A\nu\gamma$ by way of the equivalent photon approximation (EPA). The EPA treats the nucleus as a static charge distribution which sources a Coulomb field coherently (see Ref. [96, 97] for a comprehensive review). As discussed in Ref. [87, 98] the full $\sigma_{\nu A}$ cross section can be calculated from the $\sigma_{\gamma\nu}$ cross section via

$$\sigma_{\nu A} = \int_{s_{min}}^{s_{max}} {\rm d}s\ \sigma_{\gamma\nu}(s) \int_{\left(\frac{s}{2E_\nu}\right)^2}^{\infty} {\rm d}Q^2 P(s,Q^2) \qquad \text{(A15)}$$

where E_ν is the energy of the neutrino in the lab frame. The function $P(s,Q^2)$ can be interpreted as the probability of the nucleus sourcing a quasi-real photon with "mass" Q^2 whose center of mass energy with the incident neutrino is s. Typically the EPA reveals an IR logarithmic enhancement, due to the effective measure of ds/s induced by $P(s,Q^2)$. This IR enhancement is offset due to the steep s dependence of $\sigma_{\gamma\nu}(s)$ [95], which scales as $\sigma_{\gamma\nu}(s) \propto s^{2.8}$ for $s \to 0$. Using the EPA approximation to calculate production in the lead bricks of SHiP for a representative neutrino energy of $E_\nu = 20$ GeV, we find a SM background estimate of

$$\frac{\sigma_{\rm bkg}}{\text{Pb atom}} = 5.7 \cdot 10^{-10}\text{fb} = 5.7 \cdot 10^{-49}\text{cm}^2. \qquad \text{(A16)}$$

This is many orders of magnitude lower than the HNL production cross section estimated in the previous section and can safely be ignored. The smallness of this process follows physically from Yang's theorem [94, 99, 100].

Appendix B: Sensitivity

We wish to briefly outline the general strategy for how all of the projected and real exclusion limits were calculated. The strategy is based on the 2009 PDG on statistics [101]. We consider a counting experiment where the experiment has seen n events, whereas b were predicted from the Standard Model and s from new physics. In a Bayesian framework given a posterior probability and likelihood function, one can set an upper limit at credibility level $1-\alpha$ by solving

$$1-\alpha = \int_0^{s_{\rm up}} p(s|n)ds = \frac{\int_{-\infty}^{s_{\rm up}} L(n|s)\pi(s)ds}{\int_{-\infty}^{\infty} L(n|s)\pi(s)ds}. \qquad \text{(B1)}$$

Using a flat prior in the new physics signal rate and the Poisson likelihood function

$$L(n|s) = \frac{(s+b)^n}{n!}{\rm e}^{-(s+b)}, \qquad \text{(B2)}$$

this can be rewritten as

$$\begin{aligned}\alpha &= {\rm e}^{-s_{up}} \frac{\sum_{m=0}^{n}(s_{\rm up}+b)^m/m!}{\sum_{m=0}^{n} b^m/m!} \\ &= \frac{\Gamma_{\text{upper incomplete}}(1+n, b+s_{\rm up})}{\Gamma_{\text{upper incomplete}}(1+n,b)}.\end{aligned} \qquad \text{(B3)}$$

Solving for s_{up} gives us the number of signal events consistent with the observation and background prediction at $(1-\alpha)$CL. Throughout this paper, we choose $1-\alpha = 95\%$. To estimate projected sensitivities, we assume that $n = b$, namely that the data collected exactly matches the background prediction. For the LHC data, we implement the CL_s method due to the presence of under-fluctuations of the data compared to the background predictions. This consists in defining

$$\alpha_b = \int_n^{\infty} L(n'|b)dn' \qquad \text{(B4)}$$

and solving for s_{up} in

$$\alpha' \equiv \frac{\alpha}{1-\alpha_b} = 5\%, \tag{B5}$$

with α defined in Eq. (B3). This method overcovers in order to avoid setting bounds to signal rates which we are insensitive to, which can happen precisely when the data under-fluctuates. In all these cases, once we have obtained s_{up}, we can solve for the new physics coupling in the equation

$$s_{up} = \mathcal{L}\sigma_{\rm prod}\mathrm{Br}(N\to\gamma\nu)\epsilon_{\rm cuts}A_{\rm geom}P_{\rm dec}(L_1,L_2). \tag{B6}$$

In the equation above, $\mathcal{L}$ is the luminosity of the experiment. In the case of beam dump experiments, there is often an implicit sum over neutrino energies, and $\mathcal{L}$ is obtained by considering the rates and cross section of CC events in the experiment, as thoroughly described in [87].

Appendix C: Analytic cuts

The calculations applicable for the neutrino experiments and for LEP are all done analytically. We generically proceed by calculating on-shell production of the HNL in the geometric region of interest and apply efficiency cuts to ensure that the recoiling nucleus/nucleon and outgoing photon from the decay of N have the correct properties. It is thus important to devise handles that allow us to estimate these cuts as a function of the energy of N.

Consider the reaction $\nu(p_1)A(p_2)\to N(p_3)A'(p_4)$, followed by $N(p_3)\to\gamma(q_2)\nu(q_1)$. It is a relatively simple exercise in field theory to obtain $d\sigma/dt$. From here, we must determine the bounds on t. When working in the coherent scattering regime, we limit ourselves to the range $-0.5\ \mathrm{GeV}^2 < t < 0$. For inelastic scattering, we limit ourselves to $-2\ \mathrm{GeV}^2 < t < -(0.217\ \mathrm{GeV})^2/A^{\frac{2}{3}}$, and the $t < -2\ \mathrm{GeV}^2$ region applies for DIS. Within these regions, we need to pick bounds on t such that when evaluated in the lab frame, the angle of the HNL overlaps with the detector. We can further restrict the range of t by considering recoil cuts on the outgoing nucleus and nucleon respectively. Assuming p_2 initially starts at rest we have

$$\begin{aligned} t &= (p_4-p_2)^2 \\ &= 2M_H^2 - 2M_H E_4 \\ \Rightarrow t &\le 2M_H^2 - 2M_H E_4^{\text{tot cut}} \end{aligned} \tag{C1}$$

where M_H is the mass of the nucleus or nucleon depending on the context. For completeness, we also derive that

$$E_N = E_\nu + \frac{t}{2M_H}. \tag{C2}$$

This equation will be convenient for limiting ourselves to values of t in which E_N is 4 times above the photon energy threshold of the experiment, and to ensure that E_N is sufficiently boosted for production in the line of sight.

In addition, we derive cuts that require the photon from $N\to\gamma\nu$ to point in the forward direction and be above the energy threshold of the experiment. Following helicity discussions in [7, 102], we assume that the photon from the HNL decay is emitted according to a $1-\cos\vartheta$ distribution, where ϑ is the photon angle in the HNL rest frame with respect to the boost direction. We sample this distribution using an inverse transform method [103] that takes the cumulative distribution function as a random variable between 0 and 1. The energy of the photon in the rest frame is $\mathcal{E}=m_N/2$. A standard derivation shows that in the lab frame where the HNL has a boost factor γ, the photon's energy and the cosine of its angle are mapped to

$$\begin{aligned} E_{\gamma\rm lab} &= \gamma\mathcal{E}\left(1+\beta\cos\vartheta\right) \\ \cos\theta_{\rm lab} &= \frac{\gamma\left(\cos\vartheta+\beta\right)}{\sqrt{\left(\gamma\cos\vartheta+\beta\gamma\right)^2+\sin^2\vartheta}}. \end{aligned} \tag{C3}$$

We will typically choose $\theta_{\rm lab} < \pi/4$, since we want to emit photons in a cone centered along the initial direction of N. An example of the efficiency achieved using this Monte Carlo procedure is shown in Fig. 12.

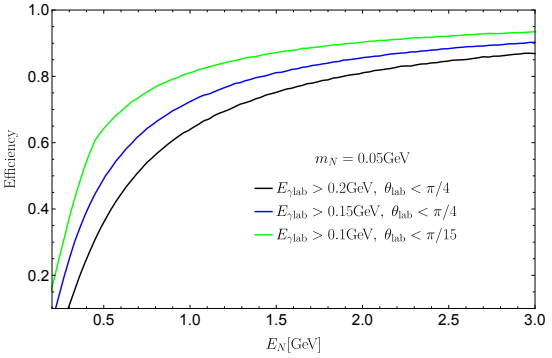

FIG. 12. Efficiency of imposing energy and angular cuts on the photon produced by a 50 MeV heavy neutral lepton.

[1] T. Asaka and M. Shaposhnikov, Phys. Lett. **B620**, 17 (2005), arXiv:hep-ph/0505013 [hep-ph].

[2] D. Gorbunov and M. Shaposhnikov, JHEP **10**, 015 (2007), [Erratum: JHEP11,101(2013)], arXiv:0705.1729

[hep-ph].
[3] S. Alekhin *et al.*, Rept. Prog. Phys. **79**, 124201 (2016), arXiv:1504.04855 [hep-ph].
[4] S. N. Gninenko, Phys. Rev. Lett. **103**, 241802 (2009), arXiv:0902.3802 [hep-ph].
[5] S. N. Gninenko, Phys. Rev. **D83**, 015015 (2011), arXiv:1009.5536 [hep-ph].
[6] D. McKeen and M. Pospelov, Phys. Rev. **D82**, 113018 (2010), arXiv:1011.3046 [hep-ph].
[7] M. Masip, P. Masjuan, and D. Meloni, JHEP **01**, 106 (2013), arXiv:1210.1519 [hep-ph].
[8] M. Masip and P. Masjuan, Phys. Rev. **D83**, 091301 (2011), arXiv:1103.0689 [hep-ph].
[9] S. N. Gninenko, Phys. Lett. **B710**, 86 (2012), arXiv:1201.5194 [hep-ph].
[10] C. Giunti and A. Studenikin, Rev. Mod. Phys. **87**, 531 (2015), arXiv:1403.6344 [hep-ph].
[11] A. Aparici, K. Kim, A. Santamaria, and J. Wudka, Phys. Rev. **D80**, 013010 (2009), arXiv:0904.3244 [hep-ph].
[12] A. Aparici, *Exotic properties of neutrinos using effective Lagrangians and specific models*, Ph.D. thesis, Valencia U. (2013), arXiv:1312.0554 [hep-ph].
[13] A. Caputo, P. Hernandez, J. Lopez-Pavon, and J. Salvado, JHEP **06**, 112 (2017), arXiv:1704.08721 [hep-ph].
[14] P. Coloma, P. A. N. Machado, I. Martinez-Soler, and I. M. Shoemaker, Phys. Rev. Lett. **119**, 201804 (2017), arXiv:1707.08573 [hep-ph].
[15] K. N. Abazajian, Phys. Rept. **711-712**, 1 (2017), arXiv:1705.01837 [hep-ph].
[16] C. Athanassopoulos *et al.* (LSND), Phys. Rev. Lett. **77**, 3082 (1996), arXiv:nucl-ex/9605003 [nucl-ex].
[17] A. A. Aguilar-Arevalo *et al.* (MiniBooNE), Phys. Rev. Lett. **98**, 231801 (2007), arXiv:0704.1500 [hep-ex].
[18] R. J. Hill, Phys. Rev. **D84**, 017501 (2011), arXiv:1002.4215 [hep-ph].
[19] K. N. Abazajian *et al.*, (2012), arXiv:1204.5379 [hep-ph].
[20] F. Couchot, S. Henrot-Versillé, O. Perdereau, S. Plaszczynski, B. Rouillé D'Orfeuil, M. Spinelli, and M. Tristram, Astron. Astrophys. (2017), 10.1051/0004-6361/201730927, [Astron. Astrophys.606,A104(2017)], arXiv:1703.10829 [astro-ph.CO].
[21] B. Dasgupta and J. Kopp, Phys. Rev. Lett. **112**, 031803 (2014), arXiv:1310.6337 [hep-ph].
[22] S. Hannestad, R. S. Hansen, and T. Tram, Phys. Rev. Lett. **112**, 031802 (2014), arXiv:1310.5926 [astro-ph.CO].
[23] M. Antonello *et al.* (LAr1-ND, ICARUS-WA104, MicroBooNE), (2015), arXiv:1503.01520 [physics.ins-det].
[24] P. Ballett, S. Pascoli, and M. Ross-Lonergan, JHEP **04**, 102 (2017), arXiv:1610.08512 [hep-ph].
[25] V. Barger, W.-Y. Keung, and D. Marfatia, Phys. Lett. **B696**, 74 (2011), arXiv:1007.4345 [hep-ph].
[26] R. Shrock, Phys. Rev. D **9**, 743 (1974).
[27] P. B. Pal and L. Wolfenstein, Phys. Rev. D **25**, 766 (1982).
[28] R. E. Shrock, Nuclear Physics B **206**, 359 (1982).
[29] R. N. Mohapatra, Phys. Rev. Lett. **56**, 561 (1986).
[30] R. N. Mohapatra and J. W. F. Valle, *Proceedings, 23RD International Conference on High Energy Physics, JULY 16-23, 1986, Berkeley, CA*, Phys. Rev. **D34**, 1642 (1986).
[31] B. Grzadkowski, M. Iskrzynski, M. Misiak, and J. Rosiek, JHEP **10**, 085 (2010), arXiv:1008.4884 [hep-ph].
[32] A. A. Aguilar-Arevalo *et al.* (MiniBooNE Collaboration), Phys. Rev. D **79**, 072002 (2009).
[33] L. B. Auerbach *et al.* (LSND), Phys. Rev. **D63**, 112001 (2001), arXiv:hep-ex/0101039 [hep-ex].
[34] E. Izaguirre, G. Krnjaic, and M. Pospelov, Phys. Lett. **B740**, 61 (2015), arXiv:1405.4864 [hep-ph].
[35] Y. Kahn, G. Krnjaic, J. Thaler, and M. Toups, Phys. Rev. **D91**, 055006 (2015), arXiv:1411.1055 [hep-ph].
[36] M. Pospelov and Y.-D. Tsai, Phys. Lett. **B785**, 288 (2018), arXiv:1706.00424 [hep-ph].
[37] R. Burman, M. Potter, and E. Smith, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **291**, 621 (1990).
[38] R. Burman and E. Smith, LAMPF Report **LA-11502-MS** (1989), 10.2172/6167579.
[39] V. Bernard, T. R. Hemmert, and U.-G. Meissner, Nucl. Phys. **A686**, 290 (2001), arXiv:nucl-th/0001052 [nucl-th].
[40] L. Alvarez-Ruso and E. Saul-Sala, in *Proceedings, Prospects in Neutrino Physics (NuPhys2016): London, UK, December 12-14, 2016* (2017) arXiv:1705.00353 [hep-ph].
[41] F. Vannucci, Adv. High Energy Phys. **2014**, 129694 (2014).
[42] J. Altegoer *et al.*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **404**, 96 (1998).
[43] J. Altegoer *et al.*, Physics Letters B **428**, 197 (1998).
[44] S. N. Gninenko and N. V. Krasnikov, Phys. Lett. **B450**, 165 (1999), arXiv:hep-ph/9808370 [hep-ph].
[45] S. Gninenko and N. Krasnikov, Physics Letters B **427**, 307 (1998).
[46] M. Anelli *et al.* (SHiP), (2015), arXiv:1504.04956 [physics.ins-det].
[47] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, JHEP **07**, 079 (2014), arXiv:1405.0301 [hep-ph].
[48] A. Alloul, N. D. Christensen, C. Degrande, C. Duhr, and B. Fuks, Comput. Phys. Commun. **185**, 2250 (2014), arXiv:1310.1921 [hep-ph].
[49] N. D. Christensen and C. Duhr, Comput. Phys. Commun. **180**, 1614 (2009), arXiv:0806.4194 [hep-ph].
[50] O. Adriani *et al.*, Phys. Lett. B **297**, 469 (1992).
[51] R. Akers *et al.* (OPAL), Z. Phys. **C65**, 47 (1995).
[52] P. Abreu *et al.* (DELPHI), Z. Phys. **C74**, 577 (1997).
[53] J. L. Lopez, D. V. Nanopoulos, and A. Zichichi, Phys. Rev. **D55**, 5813 (1997), arXiv:hep-ph/9611437 [hep-ph].
[54] R. Assmann, M. Lamont, and S. Myers, *The legacy of LEP and SLC. Proceedings, 7th Topical Seminar, Siena, Italy, October 8-11, 2001*, Nucl. Phys. Proc. Suppl. **109B**, 17 (2002).
[55] M. Aaboud *et al.* (ATLAS), Eur. Phys. J. **C77**, 393 (2017), arXiv:1704.03848 [hep-ex].
[56] M. Aaboud *et al.* (ATLAS), Eur. Phys. J. **C76**, 666 (2016), arXiv:1606.01813 [hep-ex].
[57] V. Khachatryan *et al.* (CMS), Phys. Lett. **B757**, 6 (2016), arXiv:1508.01218 [hep-ex].
[58] O. Ruchayskiy and A. Ivashko, JCAP **1210**, 014 (2012), arXiv:1202.2841 [hep-ph].

[59] S. Dodelson and L. M. Widrow, Phys. Rev. Lett. **72**, 17 (1994), arXiv:hep-ph/9303287 [hep-ph].
[60] B. D. Fields, P. Molaro, and S. Sarkar, Chin. Phys. **C38**, 339 (2014), arXiv:1412.1408 [astro-ph.CO].
[61] G. G. Raffelt, *Stars as laboratories for fundamental physics* (University of Chicago Press, 1996).
[62] H. K. Dreiner, C. Hanhart, U. Langenfeld, and D. R. Phillips, Phys. Rev. **D68**, 055004 (2003), arXiv:hep-ph/0304289 [hep-ph].
[63] H. K. Dreiner, H. E. Haber, and S. P. Martin, Phys. Rept. **494**, 1 (2010), arXiv:0812.1594 [hep-ph].
[64] H. K. Dreiner, J.-F. Fortin, C. Hanhart, and L. Ubaldi, Phys. Rev. **D89**, 105015 (2014), arXiv:1310.3826 [hep-ph].
[65] T. Fischer, S. Chakraborty, M. Giannotti, t. Mirizzi, A. Payez, and A. Ringwald, Phys. Rev. **D94**, 085012 (2016), arXiv:1605.08780 [astro-ph.HE].
[66] J. H. Chang, R. Essig, and S. D. McDermott, JHEP **01**, 107 (2017), arXiv:1611.03864 [hep-ph].
[67] E. Hardy and R. Lasenby, JHEP **02**, 033 (2017), arXiv:1611.05852 [hep-ph].
[68] K. Hirata, T. Kajita, M. Koshiba, M. Nakahata, Y. Oyama, N. Sato, A. Suzuki, M. Takita, Y. Totsuka, T. Kifune, T. Suda, K. Takahashi, T. Tanimori, K. Miyano, M. Yamada, E. W. Beier, L. R. Feldscher, S. B. Kim, A. K. Mann, F. M. Newcomer, R. Van, W. Zhang, and B. G. Cortez, Phys. Rev. Lett. **58**, 1490 (1987).
[69] E. N. Alekseev, L. N. Alekseeva, I. V. Krivosheina, and V. I. Volchenko, Phys. Lett. **B205**, 209 (1988).
[70] R. M. Bionta *et al.*, Phys. Rev. Lett. **58**, 1494 (1987).
[71] P. Gondolo and G. Gelmini, Nucl. Phys. **B360**, 145 (1991).
[72] M. Cannoni, Phys. Rev. **D89**, 103533 (2014), arXiv:1311.4494 [astro-ph.CO].
[73] H. K. Dreiner, C. Hanhart, U. Langenfeld, and D. R. Phillips, Phys. Rev. **D68**, 055004 (2003), arXiv:hep-ph/0304289 [hep-ph].
[74] A. Fradette and M. Pospelov, Phys. Rev. **D96**, 075033 (2017), arXiv:1706.01920 [hep-ph].
[75] S. A. Colgate and R. H. White, Astrophys. J. **143**, 626 (1966).
[76] M. Rampp and H. T. Janka, Astrophys. J. **539**, L33 (2000), arXiv:astro-ph/0005438 [astro-ph].
[77] M. Liebendoerfer, A. Mezzacappa, F.-K. Thielemann, O. E. B. Messer, W. R. Hix, and S. W. Bruenn, Phys. Rev. **D63**, 103004 (2001), arXiv:astro-ph/0006418 [astro-ph].
[78] B. Müller, T. Melson, A. Heger, and H. T. Janka, Mon. Not. Roy. Astron. Soc. **472**, 491 (2017), arXiv:1705.00620 [astro-ph.SR].
[79] G. T. Zatsepin and A. Yu. Smirnov, Pisma Zh. Eksp. Teor. Fiz. **28**, 379 (1978).
[80] G. M. Fuller, A. Kusenko, and K. Petraki, Phys. Lett. **B670**, 281 (2009), arXiv:0806.4273 [astro-ph].
[81] J. P. Chou, D. Curtin, and H. J. Lubatti, Phys. Lett. **B767**, 29 (2017), arXiv:1606.06298 [hep-ph].
[82] V. V. Gligorov, S. Knapen, M. Papucci, and D. J. Robinson, Phys. Rev. **D97**, 015023 (2018), arXiv:1708.09395 [hep-ph].
[83] J. Feng, I. Galon, F. Kling, and S. Trojanowski, Phys. Rev. **D97**, 035001 (2018), arXiv:1708.09389 [hep-ph].
[84] F. Kling and S. Trojanowski, (2018), arXiv:1801.08947 [hep-ph].
[85] J. C. Helo, M. Hirsch, and Z. S. Wang, (2018), arXiv:1803.02212 [hep-ph].
[86] G. Magill, R. Plestid, M. Pospelov, and Y.-D. Tsai, (2018), arXiv:1806.03310 [hep-ph].
[87] G. Magill and R. Plestid, Phys. Rev. **D95**, 073004 (2017), arXiv:1612.05642 [hep-ph].
[88] U. D. Jentschura and V. G. Serbo, Eur. Phys. J. **C64**, 309 (2009), arXiv:0908.3853 [hep-ph].
[89] C. F. Perdrisat, V. Punjabi, and M. Vanderhaeghen, Prog. Part. Nucl. Phys. **59**, 694 (2007), arXiv:hep-ph/0612014 [hep-ph].
[90] D. H. Beck and B. R. Holstein, Int. J. Mod. Phys. **E10**, 1 (2001), arXiv:hep-ph/0102053 [hep-ph].
[91] P. deNiverville, C.-Y. Chen, M. Pospelov, and A. Ritz, Phys. Rev. **D95**, 035006 (2017), arXiv:1609.01770 [hep-ph].
[92] M. Bonesini, A. Marchionni, F. Pietropaolo, and T. Tabarelli de Fatis, Eur. Phys. J. **C20**, 13 (2001), arXiv:hep-ph/0101163 [hep-ph].
[93] C. Mariani, G. Cheng, J. M. Conrad, and M. H. Shaevitz, Phys. Rev. **D84**, 114021 (2011), arXiv:1110.0417 [hep-ex].
[94] D. A. Dicus and W. W. Repko, Phys. Rev. Lett. **79**, 569 (1997), arXiv:hep-ph/9703210 [hep-ph].
[95] A. Abbasabadi, A. Devoto, and W. W. Repko, Phys. Rev. **D63**, 093001 (2001), arXiv:hep-ph/0012257 [hep-ph].
[96] R. Belusevic and J. Smith, Phys. Rev. D **37**, 2419 (1988).
[97] V. M. Budnev, I. F. Ginzburg, G. V. Meledin, and V. G. Serbo, Phys. Rept. **15**, 181 (1975).
[98] G. Magill and R. Plestid, Phys. Rev. **D97**, 055003 (2018), arXiv:1710.08431 [hep-ph].
[99] M. Gell-Mann, Phys. Rev. Lett. **6**, 70 (1961).
[100] C. N. Yang, Phys. Rev. **77**, 242 (1950).
[101] C. Amsler *et al.* (Particle Data Group), Physics Letters **B667**, 1 (2008 and 2009 partial update for the 2010 edition).
[102] L. F. Li and F. Wilczek, Phys. Rev. D **25**, 143 (1982).
[103] C. Patrignani *et al.* (Particle Data Group), Chin. Phys. **C40**, 100001 (2016).

Chapter 5

Neutrino trident production at the intensity frontier

5.1 Preface

In this paper, we calculate rates at the future planned experiments of DUNE and SHiP for a process called neutrino trident production (NTP). In this process, an incoming neutrino scatters with a photon sourced by a nucleus to produce two oppositely charged leptons and an outgoing neutrino. This process has mainly been considered (and observed) in the context of two oppositely charged muons. We calculate rates at future experiments for this process to occur for an oppositely charged electron/muon pair, as well as an electron/positron pair. These rates don't vanish within the SM model, and for some modes even surpass by an order of magnitude those of the conventional muon channel. We also consider final states involving tauons, which we find to be very small. Our paper was the first to point out that SM induced mixed-flavor trident processes is observable at upcoming experiments.

In our calculations, we employ the equivalent photon approximation (EPA). This approximation is valid for large incoming neutrino energies, and for kinematics in which the photon sourced by the nucleus is real and on-shell. In this approximation, the momentum of the photon, q, satisfies $q^2 = 0$ and parts of the

phase space kinematics can be simplified, essentially reducing to a 3-body phase space. Through discussions with colleagues that are doing follow-up work on mixed-flavor trident which is currently in preparation, it turns out that comparing this method to a full 4-body phase space treatment, the EPA overestimates the cross section by a factor of 3-4 at DUNE for mixed-flavor final states. At the much larger incoming neutrino energies of SHiP, the discrepancy is much smaller since the approximation is better. For heavier final states such as two oppositely charged muons or final states involving tauons, the EPA and the full 4-body treatment are in agreement.

This said, the estimates we've made at DUNE were based on initial designs of a 100kg near detector. Recently, the near detector designs were upgraded to a 30 tonne near detector. Since the rates for trident scale approximately linearly in the mass of the detector, this increases our rates by a factor of 300. As such, mixed-flavor NTP will still be very much observable at DUNE. Partially as a result of our paper, this process is receiving much closer attention by the particle physics community, and is currently being implemented in the full simulations and physics programme of the DUNE collaboration.

 This paper appears in https://doi.org/10.1103/PhysRevD.95.073004 and is published in PRD as

- Gabriel Magill and Ryan Plestid, "Neutrino trident production at the intensity frontier," Phys. Rev. **D95**, 073004 (2017), arXiv:1612.05642 [hep-ph].

Contribution of author (reproduced from Declaration of Authorship): Both authors contributed significantly to the calculations, writing and editing of the paper.

5.2 Paper

PHYSICAL REVIEW D **95,** 073004 (2017)

Neutrino trident production at the intensity frontier

Gabriel Magill[1,2,*] and Ryan Plestid[1,2,†]

[1]*Department of Physics & Astronomy, McMaster University, 1280 Main Street West, Hamilton, L8S 4M1 Ontario, Canada*

[2]*Perimeter Institute for Theoretical Physics, 31 Caroline Street, North, Waterloo, N2L 2Y5 Ontario, Canada*

(Received 1 March 2017; published 13 April 2017)

We have calculated cross sections for the production of lepton pairs by a neutrino incident on a nucleus using both the equivalent photon approximation and deep-inelastic formalism. We find that production of mixed flavor lepton pairs can have production cross sections as high as 35 times those of the traditional $\nu_\mu \to \nu_\mu \mu^+ \mu^-$ process. Rates are estimated for the SHiP and DUNE intensity frontier experiments. We find that multiple trident production modes, some of which have never been observed, represent observable signals over the lifetime of the detectors. Our estimates indicate that the SHiP Collaboration should be able to observe on the order of 300 trident events given 2×10^{20} protons on target and that the DUNE Collaboration can expect approximately 250 trident events in their near detector given 3×10^{22} protons on target. We also discuss possible applications of the neutrino trident data to be collected at SHiP and DUNE for Standard Model and beyond the Standard model physics.

DOI: 10.1103/PhysRevD.95.073004

I. INTRODUCTION

Neutrino physics has traditionally been dominated by the measurement of oscillation parameters and the study of neutrino nucleus scattering. These experimental signals are largely dominated by charged-current (CC) and neutral-current (NC) interactions of which the cross sections scale as $\sigma \sim sG_F^2$. Traditionally, limits on beam luminosity have resulted in event counts that leave subdominant processes with expected event rates less than unity in the lifetime of an experiment. As a result, these processes are often omitted in the discussions of neutrino physics. One such neglected process is neutrino trident production which has been previously observed at CHARM II, CCFR, and NuTev [1–3]. These measurements provided evidence at the 3σ level for the contribution of Z bosons in weak interactions [2] and more recently have been used to constrain beyond the Standard Model (BSM) physics. Specifically, measurements from CCFR currently provide the best constraints on the mass and coupling of a heavy Z' force mediator charged under $L_\mu - L_\tau$ [4]. Both of these applications are successful because the neutrino trident production of leptons is sensitive to both the vector- and axial-current couplings (see Sec. II A).

The aforementioned collaborations only measured one possible mode of trident production, specifically $\nu A \to \nu \mu^+ \mu^- A$. The leading order contribution to this process involves the production of a muon-antimuon pair, which can then interact with the target nucleus A electromagnetically (see Fig. 1). For low momentum transfers ($Q \ll R_A^{-1}$), the nucleus interacts coherently with the virtual photons ($\sigma \propto Z^2$), and there is a strong enhancement due to the infrared divergence in the photon propagator; it is this kinematic regime which dominates the cross section. Other qualitatively similar processes, such as e^+e^- or μ^+e^- trident production, were kinematically accessible; however, due to technological limitations in the detector design, the required vertex resolution for trident identification was not achievable for electrons. This will not be an issue with modern detectors.

The cross section for $\mu^+\mu^-$ neutrino trident production is approximately 5 orders of magnitude smaller than the charged-current cross section ($\sigma \approx 10^{-5}\sigma_{CC}$) for a 50 GeV neutrino scattering off an iron nucleus [5]; high Z materials will have an even larger cross section relative to CC scattering. This means that practically trident production can only be observed in experiments with very large neutrino fluxes. Additionally, the leading contribution to the cross section discussed in the preceding paragraph can be calculated using the equivalent photon approximation and scales as $\sigma \sim G_F^2 E_\nu Q_{\max} \log(E_\nu Q_{\max}/m_\ell^2)$, where m_ℓ is related to the lepton masses and $Q_{\max}$ is a characteristic

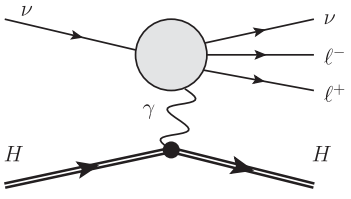

FIG. 1. Leading hadronic contribution to trident production. Arrows denote the direction of momentum.

*gmagill@perimeterinstitute.ca
†plestird@mcmaster.ca

momentum transfer set by the radius of the nucleus [5]. These considerations imply that for trident production to be a useful tool one needs to consider experiments with both a high energy neutrino beam ($\langle E_\nu \rangle \gtrsim 1$ GeV) and high statistics. This can be achieved via beam luminosity, or target-mass considerations. Fixed target and beam dump experiments—where neutrino energies can be in excess of 100 GeV, and charged-current event counts can exceed 10^6—are an ideal setting to study neutrino trident production. The SHiP experiment [6] and DUNE [7] both fall into these categories and, as we show in this paper, represent the newest frontier in the study of trident production.

SHiP's program of study, as it relates to neutrino physics, is largely focused on tau neutrino and antitau neutrino events and is therefore optimized to observe tau leptons [6]. This represents a qualitatively new opportunity in the study of trident production, because the high mass of the tau leptons results in a threshold effect, wherein coherent production of a single tau lepton is not possible unless the inequality $E_\nu > (1/2)m_\tau^2 R_A$ holds; the bound for tau lepton pair production is given by $E_\nu > 2m_\tau^2 R_A$. As a result, we also investigate the incoherent contribution to the cross section using both a diffractive and deep-inelastic approach. The experiment will use beams with $\langle E_\nu \rangle \approx 30$ GeV $-$ 60 GeV and expects a lifetime collection of charged-current events on the order of $N_{\mathrm{CC}} \approx 2.7 \times 10^6$ [6]. It is therefore reasonable to assume that mixed flavor trident production, possibly including tau leptons, should be observable at the SHiP experiment.

Although the focus of its program of study is neutrino oscillations, DUNE will use sufficiently high luminosities and neutrino energies to induce trident production. The experiment consists of a near detector on site at FERMILAB [7] and a far detector at Sanford Laboratory, both composed of liquid argon. This technology allows for the observation of both electrons and muons. The far detector is exposed to a flux of neutrinos after a 1300 km transit through Earth. The near detector will be used to account for systematic uncertainties in the neutrino beam and to record the initial neutrino flux. It is designed to obtain ten times the statistics of the far detector [7]. The expected charged-current event count in the far detector over the lifetime of the experiment is on the order of 1×10^5, and so it is reasonable to expect an observable signal of trident events for some of the processes, especially given the enhanced statistics of the planned near detector.

Trident production has proven itself a useful tool for constraining BSM physics by virtue of its sensitivity to modifications of C_A and C_V. Additionally, it represents an experimental signal that would provide an obvious background to searches of lepton flavor violation in the case of multiflavor charged-lepton tridents. If these new experiments (SHiP and DUNE) were to use trident production to probe BSM physics, then it would be imperative to understand the relevant Standard Model backgrounds.

The rest of this article is organized as follows. In Sec. II A, we discuss the basic structure of the trident amplitude in the Standard Model. In Sec. II B, we describe how to obtain the cross sections for three distinct kinematic regimes, each receiving a separate theoretical treatment. In Sec. III, we calculate expected rates and cross sections for both DUNE and SHiP. We also present differential distributions with respect to the invariant mass of the charged-lepton pair. In Sec. IV, we review the qualitative features of our results and outline possible applications of trident production for both SHiP and DUNE. Finally, in Sec. V, we discuss future directions for trident production for the upcoming generation of accelerator-based neutrino experiments.

II. TRIDENT PRODUCTION IN THE STANDARD MODEL

A. Leptonic matrix element

Our treatment of trident production varies over kinematic regimes, characterized by the four-momentum transfer to the nucleus Q^2. In every approach, we treat the leptonic matrix element involving the electromagnetic current consistently. Our treatment of the nucleus's interaction with the electromagnetic field, however, varies and so will be treated separately in each section. In the lower Q^2 regimes, we relate the cross section to that of a neutrino-photon collision (photo-trident production), while for large Q^2, we employ the parton model. The amplitudes for photo-trident production and parton-trident production can be written

$$\begin{aligned} \imath\mathcal{M}_{\gamma\nu} &= \epsilon^\mu L_\mu \qquad \text{(equivalent photon approximation)} \\ \imath\mathcal{M}_{h\nu} &= \frac{-\eta^{\mu\nu}}{q^2} h_\nu L_\mu \quad \text{(deep-inelastic scattering)} \end{aligned} \tag{1}$$

where ϵ^μ is an on-shell polarization tensor and h_ν is the hadronic matrix element in the parton model. The leptonic matrix element L_μ is calculated explicitly below. We study both neutrino and antineutrino induced trident production, and for the remainder of this section, all reactions will contain an implicit hadronic initial and final state. We use latin flavor indices $i, j, k \in \{e, \mu, \tau\}$ and consider reactions of the form

$$\{\nu_i \to \nu_{i \text{ or } k} + \ell_j^- + \ell_k^+, \bar{\nu}_i \to \bar{\nu}_{i \text{ or } j} + \ell_j^- + \ell_k^+\}$$

with the constraint that the generational lepton number is conserved. Both monoflavor and multiflavor charged-lepton pairs (i.e. $\mu^+\mu^-$ and $\mu^+\tau^-$) are included in our analysis. Assigning the labels $\{1, 2, 3, 4, 5\} \to \{\nu, \gamma, \nu', \ell^+, \ell^-\}$ with ν' the outgoing neutrino (see Fig. 2) and generalizing the analysis of Refs. [4,8] to multiflavor lepton pairs, we find

$$\begin{aligned} L^\mu_{ijk} = &-\frac{\imath e G_F}{\sqrt{2}} \{\bar{u}_3 \gamma^\alpha (1-\gamma^5) u_1, \bar{v}_1 \gamma^\alpha (1-\gamma^5) v_3\} \\ &\times \bar{u}_5 \left[\gamma_\alpha (V_{ijk} - A_{ijk}\gamma^5) \frac{1}{\not{q} - \not{p}_4 - m_4} \gamma^\mu \right. \\ &\left. + \gamma^\mu \frac{1}{\not{p}_5 - \not{q} - m_5} \gamma_\alpha (V_{ijk} - A_{ijk}\gamma^5) \right] v(p_4), \end{aligned} \tag{2}$$

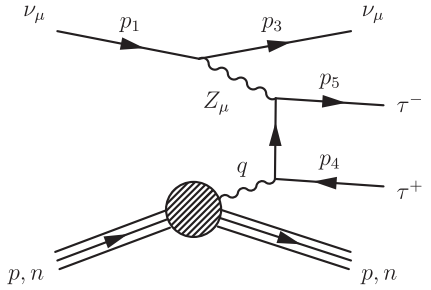

FIG. 2. An example of a process which takes place exclusively through the neutral-current channel. The mismatch in flavor between the incident neutrino and outgoing leptons prohibits a charged-current interaction.

where the first line contains the appropriate spinor wave functions for an incident neutrino and antineutrino beam respectively. V_{ijk} and A_{ijk} are the flavor dependent vector and axial coupling strengths, which are typically denoted C_V and C_A respectively. We use nonstandard notation to stress that these couplings carry flavor indices because some processes are mediated exclusively by W bosons, others exclusively by Z bosons, and some a mixture of the two. As we see from Fig. 2, these mediators modify the coupling to the vector and axial currents, as can be verified by the use of Fierz identities. As noted in Ref. [5], the interference between the neutral- and charged-current channels in the Standard Model results in a 40% reduction in the cross section compared to the $V-A$ theory prediction. Thus, by considering different combinations of leptons in the final state, the cross section can be enhanced, or suppressed, significantly. The constants A_{ijk} and V_{ijk} are presented in Table I for $\nu_\mu \to \nu_\mu \tau^+ \tau^-$ and for all trident processes with lifetime event counts greater than 0.01 at either SHiP or DUNE.

TABLE I. Modified vector and axial coupling constants for different combinations of incident neutrino flavors and final states.

ν process	$\bar{\nu}$ process	V_{ijk}	A_{ijk}	Mediator
$\nu_e \to \nu_e e^+ e^-$	$\bar{\nu}_e \to \bar{\nu}_e e^+ e^-$	$\frac{1}{2} + 2\sin^2\theta_w$	$\frac{1}{2}$	W, Z
$\nu_\mu \to \nu_\mu \mu^+ \mu^-$	$\bar{\nu}_\mu \to \bar{\nu}_\mu \mu^+ \mu^-$	$\frac{1}{2} + 2\sin^2\theta_w$	$\frac{1}{2}$	W, Z
$\nu_e \to \nu_\mu \mu^+ e^-$	$\bar{\nu}_e \to \bar{\nu}_\mu e^+ \mu^-$	1	1	W
$\nu_\mu \to \nu_e e^+ \mu^-$	$\bar{\nu}_\mu \to \bar{\nu}_e \mu^+ e^-$	1	1	W
$\nu_e \to \nu_e \mu^+ \mu^-$	$\bar{\nu}_e \to \bar{\nu}_e \mu^+ \mu^-$	$-\frac{1}{2} + 2\sin^2\theta_w$	$-\frac{1}{2}$	Z
$\nu_\mu \to \nu_\mu e^+ e^-$	$\bar{\nu}_\mu \to \bar{\nu}_\mu e^+ e^-$	$-\frac{1}{2} + 2\sin^2\theta_w$	$-\frac{1}{2}$	Z
$\nu_\mu \to \nu_\mu \tau^+ \tau^-$	$\bar{\nu}_\mu \to \bar{\nu}_\mu \tau^- \tau^+$	$-\frac{1}{2} + 2\sin^2\theta_w$	$-\frac{1}{2}$	Z
$\nu_\mu \to \nu_\tau \mu^- \tau^+$	$\bar{\nu}_\mu \to \bar{\nu}_\tau \mu^+ \tau^-$	1	1	W
$\nu_\tau \to \nu_\mu \tau^- \mu^+$	$\bar{\nu}_\tau \to \bar{\nu}_\mu \tau^+ \mu^-$	1	1	W
$\nu_\tau \to \nu_\tau \mu^+ \mu^-$	$\bar{\nu}_\tau \to \bar{\nu}_\tau \mu^- \mu^+$	$-\frac{1}{2} + 2\sin^2\theta_w$	$-\frac{1}{2}$	Z
$\nu_\tau \to \nu_\tau e^+ e^-$	$\bar{\nu}_\tau \to \bar{\nu}_\tau e^- e^+$	$-\frac{1}{2} + 2\sin^2\theta_w$	$-\frac{1}{2}$	Z

B. Coherent, diffractive, and deep-inelastic regimes

We will begin by reviewing conventional scattering of neutrinos off of nuclei to emphasize the qualitative differences in trident production. Neutrino-nucleus scattering is dominated by charged-current events, which can be loosely partitioned into three classes for $E_\nu \gtrsim 100$ MeV: quasielastic scattering, hadronic resonance production, and deep-inelastic scattering [9]. It is only at low center-of-mass energies $E \lesssim 50$ MeV that coherent scattering via the neutral current is possible such that the reaction's cross section scales as $\sigma \sim (A-Z)^2 E_\nu^2$ with $A-Z$ the number of neutrons. In this energy regime, coherent scattering cross sections can be as much as 3 orders of magnitude larger than that predicted by a naive sum of the nucleon cross sections [10].

This limited kinematic window stands in sharp contrast to trident production where coherent contributions are possible at all energies, because the reaction is not $2 \to 2$ and the phase space is therefore less kinematically constrained. This scattering is mediated electromagnetically, and, in addition to the coherent Z^2 amplification, the photon's propagator introduces an infrared divergence that further enhances the amplitude. As is the case for coherent neutrino scattering, this regime is characterized by small momentum transfers ($Q^2 \sim R_A^{-2}$) wherein the phases of the various amplitudes are nearly commensurate, and the amplitudes interfere constructively. Kinematic considerations constrain the momentum transfer via $Q > s/(2E_\nu)$, with s the invariant mass of the neutrino-photon pair [5]. When combined with the lepton pair's mass threshold, this regulates the infrared divergence mentioned above. The three regimes typically considered in charged-current scattering for high energy neutrinos (mentioned in the first paragraph) also exist for trident production. Quasi-elastic-like diffractive scattering can contribute significantly to trident production, especially when threshold effects related to lepton masses are important. We expect the deep-inelastic contribution to be suppressed, but for many of the neutrino energies at SHiP, it is the only kinematically allowed production mechanism for tau leptons, and so we also include this regime in our analysis.

1. Coherent regime

The coherent contribution to neutrino trident production can be accurately calculated using the equivalent photon approximation (EPA) [4,5,11,12]. In the EPA, the cross section for the full scattering process is decomposed into two pieces. First, the cross section corresponding to the scattering of a neutrino and photon creating a lepton trident, denoted by $\sigma_{\gamma\nu}$, is calculated. Next, this cross section is weighted against a universal probability distribution $P(s, Q^2)$ [4] that measures the likelihood of the nucleus producing a virtual photon with virtual-mass Q^2 and neutrino-photon center-of-mass energy s. The full cross section is given by

$$\sigma_{\nu A} = \int \mathrm{d}s \sigma_{\gamma\nu}(s) \int \mathrm{d}Q^2 P(s, Q^2) = \frac{Z^2\alpha}{\pi} \int_{m_{jk}^2}^{s_{\max}} \frac{\mathrm{d}s}{s} \sigma_{\gamma\nu}(s) \int_{(s/2E_\nu)^2}^{\infty} \frac{\mathrm{d}Q^2}{Q^2} F^2(Q^2) \quad (3)$$

with $m_{jk} = m_j + m_k$ the sum of the lepton pair's masses. A fairly good, albeit crude, approximation is to treat the form factor for the nucleus $F(Q^2)$ as a Heaviside function $\Theta(Q^2_{\max} - Q^2)$ where the scale $Q_{\max} = \Lambda_{\mathrm{QCD}}/A^{1/3}$ corresponds to characteristic momentum transfer at which one would expect the dissolution of the nucleus [5]. This sets a maximum center-of-mass energy for the photon-neutrino interaction $s_{\max} = 2E_\nu Q_{\max}$. With these approximations, suppressing flavor indices and working in the leading log approximation, Eq. (3) simplifies to [4,5]

$$\sigma_{\nu A} \approx \frac{1}{2}(A^2 + V^2) \frac{2Z^2\alpha^2 G_F^2}{9\pi^3} s_{\max} \log\left(\frac{s_{\max}}{4m^2}\right), \quad (4)$$

where $2m = m_j + m_k$. There are additional terms resulting from the interference between the vector and axial currents, but these are suppressed by two powers of the lepton mass and are therefore small. A more realistic implementation is to use the Woods-Saxon form factor, which is what we used in all of our calculations (this changes the answer by order 10%; see Appendix A for details). We can write the coherent contribution to the neutrino-nucleus cross section as

$$\mathrm{d}\sigma_{\gamma\nu} = \frac{1}{2s}\frac{1}{2}\sum_{\mathrm{pol}} |\epsilon_\mu L^\mu|^2 \mathrm{d}\Phi_3, \quad (5)$$

where Φ_3 is the three-body phase space of final states, the factor of $1/2$ averages over photon polarizations, and $2s$ is the Lorentz-invariant flux factor. For details on the treatment of the three-body phase space, see Appendix A.

2. Diffractive regime

At intermediate Q^2, it is possible to interact with the individual protons of the nucleus, both without coherent interference of their individual amplitudes and without probing their inner parton structure. Our treatment of this regime follows the approach outlined in Ref. [13] and is identical to the coherent regime with the following changes:

$$\sigma_{\nu A} = Z \int \mathrm{d}s \sigma_{\gamma\nu}(s) \int \mathrm{d}Q^2 P(s, Q^2) = Z\frac{\alpha}{\pi} \int_{m_{jk}^2}^{s_{\max}} \frac{\mathrm{d}s}{s} \sigma_{\gamma\nu}(s) \int_{Q^2_{\min}}^{1\ \mathrm{GeV}^2} \frac{\mathrm{d}Q^2}{Q^2} F^2_{\mathrm{dip}}(Q^2). \quad (6)$$

The charge of the nucleus now appears as an overall multiplicative factor as opposed to appearing in $P(s, Q^2)$, we cut off our integral at $Q_{\min} = \max(s/2E_\nu, R_A^{-1})$ to avoid double-counting amplitudes included in the coherent calculation, and we use the standard dipole fit to the proton's electromagnetic form factor (see Appendix A). We introduce an explicit UV cutoff for the Q^2 integration to avoid double-counting with the deep-inelastic scattering (DIS) amplitudes. This was not necessary for the coherent regime due to the exponential, as opposed to power law, decay of the Wood-Saxon form factor at high Q^2.

3. Deep-inelastic regime

Our treatment of the deep-inelastic case is fairly standard, with a few exceptions that are highlighted in Appendix B 2. We treat this regime by convoluting the parton cross sections with nucleon parton distribution functions (PDFs) $f(\xi, Q)$ [14], taking into account the u, d, c, s quarks. The phase space integrals are sensitive to the lepton masses, and so, although their effects on the matrix element are often subleading, we include their full dependence throughout our calculations. All of the quarks are treated as massless in our analysis.

We take care to include a cut on momentum transfers so as not to double-count contributions already accounted for by the EPA. Additionally, we place a cut on the momentum fraction ξ to ensure the parton carries enough four-momentum to both be able to produce the appropriate pair of charged leptons and to satisfy the double-counting cut on momentum transfer. The resulting cross sections for the various nucleons are then summed to obtain the scattering cross section with the nucleus. We can write $\sigma_{\nu A}$ as a weighted sum of the cross sections with the constituent nucleons

$$\sigma_{\nu A} = Z\sigma_{\nu p} + (A - Z)\sigma_{\nu n}. \quad (7)$$

These can in turn be written in terms of the parton-level cross sections $\sigma_{h\nu}$ via

$$\sigma_{\nu H} = \sum_h \int_{\xi_{\min}}^1 \mathrm{d}\xi \int_{Q_{\min}}^{Q_{\max}} \mathrm{d}Q \frac{\mathrm{d}\sigma_{h\nu}}{\mathrm{d}Q}(\xi, Q) f_h^{(H)}(\xi, Q), \quad (8)$$

where $f_h^{(H)}(\xi, Q)$ is the PDF for parton h in the nucleon $H \in \{n, p\}$. More details can be found in Appendix B.

III. PROSPECTS AT FUTURE EXPERIMENTS

In the following, we calculate trident rates at SHiP and at the DUNE far and near detectors. We calculate the rates for the momentum transfers $Q < 0.217/(A)^{\frac{1}{3}}$ GeV $\approx R_A^{-1}$ regime using the coherent EPA method. For intermediate momentum transfers $0.217/(A)^{\frac{1}{3}}$ GeV $\lesssim Q \lesssim M_p$, we use the diffractive EPA treatment. Finally, for $Q \gtrsim 1$ GeV $\approx M_p$, we employ the deep-inelastic formalism. We use PDFs from the MSTW Collaboration (2008 next-to-next-to-leading order best fit) [14]. We use the symbolic manipulation program FORM to automate the calculation of the

spin-averaged matrix element squared [15]. To calculate the rates, we estimate the number of Standard Model (SM) neutrino trident events for each flavor of incident neutrino ν_i producing a lepton pair composed of j^- and k^+ with $i,j,k \in \{e,\mu,\tau\}$. We estimate the luminosity in terms of charged-current events N^i_{CC} using

$$N^{ijk}_{\text{Trident}} = \sum_E \frac{N^i_{CC}(E)}{\sigma_{CC}(E,A)} \sigma^{ijk}_{\nu A}(E,Z,A) \times \epsilon^j_- \times \epsilon^k_+, \quad (9)$$

where σ_{CC} is the neutrino charged-current cross sections [16] and i, j, and k are flavors denoting the incident neutrino, outgoing ℓ^-, and outgoing ℓ^+ respectively. Additionally, ϵ_+ and ϵ_- are the identification efficiencies for ℓ^+ and ℓ^- respectively. We do an analogous procedure for antineutrinos.

There will be a background contribution to trident production from resonant production of charged pions and charm production from D mesons, the leptonic modes of which are both dominated by muon flavored final states. In the different flavor opposite sign dilepton final states, backgrounds can arise from $\bar{\nu}_\mu$ CC scattering in combination with an elastic NC event releasing an electron and also by muon final states in which one of the muons fakes an electron. As coherent scattering is quasielastic, the backgrounds for the dominant contribution to the cross section (see Sec. II B) can be greatly reduced by imposing hadronic vetoes in the analysis. Further background suppression can be achieved by selecting oppositely charged leptons that fall within the vertex resolution of the detectors and selecting events with low-$M_{\ell^+\ell^-}$-invariant masses. We leave the background estimates to the collaborations' detailed and sophisticated simulations. Our signal results are shown in Tables II to IV.

A. Calibrations and tests

The details of our calculations can be found in the Appendixes. We calibrated our EPA cross section calculations with previous theoretical and experimental work [1,4,17] and reproduced the analytic results of Ref. [4].

Our DIS work was calibrated with MadGraph5 [18] for trident induced muon pair production. MadGraph5 treats light leptons as massless, and due to infrared singularities in the propagators, this necessitates a careful treatment; it also introduces questions of reliability. We imposed the following cuts to replicate the effects of finite muon masses: $p_T > m_\mu$ for the muons, $p_T > 1.5$ GeV for the jets, and $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2} > 0.4$ for the lepton pairs. With these cuts, we found our calculations to agree with MadGraph5 to within a factor of 0.5–2.5 for $E_\nu = \{20\text{ GeV}, 200\text{ GeV}, 1000\text{ GeV}\}$. We believe our calculation to be more reliable than MadGraph5 in the low-Q^2 regions of phase space which dominate the cross sections due to infrared divergences, which we treat carefully.

TABLE II. Number of expected trident events for coherent (Coh) and diffractive (Diff) scattering, using the EPA, in the SHiP ν_τ detector, assuming 2×10^{20} POT on molybdenum.

Neutrino beam			Antineutrino beam		
Process	Coh	Diff	Process	Coh	Diff
$\nu_\mu \to \nu_e e^+ \mu^-$	85.46	24.6	$\bar{\nu}_\mu \to \bar{\nu}_e e^- \mu^+$	29.96	9.61
$\nu_\mu \to \nu_\mu e^+ e^-$	28.28	5.32	$\bar{\nu}_\mu \to \bar{\nu}_\mu e^+ e^-$	22.48	3.58
$\nu_e \to \nu_e e^+ e^-$	21.69	2.95	$\bar{\nu}_e \to \bar{\nu}_e e^+ e^-$	15.65	2.45
$\nu_e \to \nu_\mu \mu^+ e^-$	9.1	2.31	$\bar{\nu}_e \to \bar{\nu}_\mu \mu^- e^+$	14.31	3.16
$\nu_\mu \to \nu_\mu \mu^+ \mu^-$	4.79	3.01	$\bar{\nu}_\mu \to \bar{\nu}_\mu \mu^+ \mu^-$	3.76	2.38
$\nu_e \to \nu_e \mu^+ \mu^-$	0.42	0.16	$\bar{\nu}_e \to \bar{\nu}_e \mu^+ \mu^-$	0.3	0.12
$\nu_\tau \to \nu_\tau e^+ e^-$	0.13	0.03	$\bar{\nu}_\tau \to \bar{\nu}_\tau e^+ e^-$	0.13	0.02
$\nu_\tau \to \nu_\tau \mu^+ \mu^-$	0.01	0.	$\bar{\nu}_\tau \to \bar{\nu}_\tau \mu^+ \mu^-$	0.01	0.
$\nu_\tau \to \tau^- \mu^+ \nu_\mu$	0.	0.01	$\bar{\nu}_\tau \to \tau^+ \mu^- \bar{\nu}_\mu$	0.	0.
$\nu_\mu \to \mu^- \tau^+ \nu_\tau$	0.	0.23	$\bar{\nu}_\mu \to \mu^+ \tau^- \bar{\nu}_\tau$	0.	0.39
Total	149.88	38.62		86.6	21.71

TABLE III. Number of expected trident events for coherent (Coh) and diffractive (Diff) scattering, using the EPA, in the lifetime of the DUNE near detector assuming $\sim 3 \times 10^{22}$ POT (equivalently, an 850 kt-MW-yr exposure at the far detector).

Neutrino beam			Antineutrino beam		
Process	Coh	Diff	Process	Coh	Diff
$\nu_\mu \to \nu_e e^+ \mu^-$	73.98	53.15	$\bar{\nu}_\mu \to \bar{\nu}_e e^- \mu^+$	25.23	18.7
$\nu_\mu \to \nu_\mu e^+ e^-$	23.03	9.64	$\bar{\nu}_\mu \to \bar{\nu}_\mu e^+ e^-$	16.45	6.79
$\nu_\mu \to \nu_\mu \mu^+ \mu^-$	2.03	5.28	$\bar{\nu}_\mu \to \bar{\nu}_\mu \mu^+ \mu^-$	2.16	4.3
$\nu_e \to \nu_e e^+ e^-$	0.7	0.29	$\bar{\nu}_e \to \bar{\nu}_e e^+ e^-$	0.54	0.22
$\nu_e \to \nu_\mu \mu^+ e^-$	0.21	0.17	$\bar{\nu}_e \to \bar{\nu}_\mu \mu^- e^+$	0.4	0.27
$\nu_e \to \nu_e \mu^+ \mu^-$	0.01	0.01	$\bar{\nu}_e \to \bar{\nu}_e \mu^+ \mu^-$	0.	0.01
Total	99.96	68.54		44.78	30.29

TABLE IV. Number of expected trident events for coherent (Coh) and diffractive (Diff) scattering, using the EPA, in the lifetime of the DUNE far detector assuming $\sim 3 \times 10^{22}$ POT (equivalently, an 850 kt-MW-yr exposure at the far detector).

Neutrino beam			Antineutrino beam		
Process	Coh	Diff	Process	Coh	Diff
$\nu_\mu \to \nu_e e^+ \mu^-$	2.12	1.52	$\bar{\nu}_\mu \to \bar{\nu}_e e^- \mu^+$	0.05	0.03
$\nu_\mu \to \nu_\mu e^+ e^-$	0.66	0.28	$\bar{\nu}_\mu \to \bar{\nu}_\mu e^+ e^-$	0.03	0.01
$\nu_e \to \nu_e e^+ e^-$	0.11	0.05	$\bar{\nu}_e \to \bar{\nu}_e e^+ e^-$	0.05	0.02
$\nu_\mu \to \nu_\mu \mu^+ \mu^-$	0.06	0.15	$\bar{\nu}_\mu \to \bar{\nu}_\mu \mu^+ \mu^-$	0.	0.01
$\nu_e \to \nu_\mu \mu^+ e^-$	0.03	0.03	$\bar{\nu}_e \to \bar{\nu}_\mu \mu^- e^+$	0.03	0.02
Total	2.98	2.03		0.16	0.09

B. Rates for SHiP

SHiP will be a lead-based neutrino detector [6,19]. It will utilize an emulsion cloud chamber for its electron detection

and a muon magnetic spectrometer for muons. It is estimated to have a 90% e and μ identification efficiency and a micron vertex resolution. Under nominal operating conditions, after five years of operation, it will have collected data from 2×10^{20} protons on target (POT) using a 400 GeV super proton synchrotron (SPS) proton beam. We quote all the rates assuming this normalization.

The energy spectrum at SHiP is very broad and reaches sufficiently high energies such that trident production of tau leptons becomes kinematically allowed in the coherent, diffractive, and deep-inelastic regimes. The latter is allowed at almost all incident neutrino energies available at SHiP with the only requirement being that the center-of-mass energy exceeds the lepton pair's mass gap. Despite being kinematically allowed, we find the large momentum transfer in the deep-inelastic regime renders the contribution to the cross section negligible. The diffractive and coherent regimes rely on the high energy tail of the quoted beam distribution [6]. For electrons and muons, coherent and diffractive production are not only possible but extremely viable, while for tau leptons, we find only diffractive production to be viable, but only marginally so. In Figs. 3 and 7, we show the cross section per nucleon as a function of the incoming neutrino energy for a variety of processes. The coherent cross sections computed via the EPA are normalized by Z^2, while the deep-inelastic contribution is normalized by A. There are small differences in these plots for various materials, as the EPA Woods-Saxon form factor and the relative number of protons to neutrons in DIS both introduce a subleading dependence on the ratio of protons to neutrons that is not removed by the per-nucleon normalization.

In Table II, we show the expected number of events in the various production modes for both low-Q^2 events calculated within the coherent EPA and intermediate-Q^2 events calculated using the diffractive EPA. DIS rates are not included, because the cumulative lifetime event count for all production modes in the deep-inelastic regime is $N_{\mathrm{DIS}}^{(\mathrm{tot})} \approx 0.1$.

The basic features of our analysis can be understood by looking at Table I and Eq. (4) and remembering that the neutrino beam is dominated by ν_μ and $\bar{\nu}_\mu$. This is discussed in greater detail in Sec. IV.

C. Rates for DUNE

DUNE [7] is composed of a near detector that primarily sees a flux of muon neutrinos and a far detector used to study the appearance of electron neutrinos as a result of oscillations from the muon neutrino beam. That said, there will be a mixture of both neutrino flavors at each site relevant for trident production. Both near and far detectors are based on argon time projection chambers, which allow for the differentiation of electrons and photons. We take the electron and muon identification efficiencies to be 90%.

In Tables III and IV, we show the expected number of events for the near and far detectors respectively. The rates in both tables are calculated assuming an 850 kt-MW-yr exposure in the far detector. This number corresponds to the amount of data collected in the lifetime of DUNE given their optimized design. To convert this measure to protons on target, note that the far detector weighs 40 kt and a beam power of 1.07 MW with 80 GeV protons corresponds to 1.47×10^{21} POT/yr [7]. This gives roughly 3×10^{22} POT. The full details of the luminosity calculations are given in Appendix C. As we did for SHiP, we consider both low-Q^2 events calculated within the coherent EPA and intermediate-Q^2 events calculated using the diffractive EPA. DIS rates are not included as they are negligible. In Fig. 4, we show the cross section per nucleon as a function of the incoming neutrino energy for each process listed in Table I, for coherent EPA. Compared to Fig. 3 there are small differences, which are due to the Woods-Saxon form factor's implicit dependence on A [see Eq. (A8) for details].

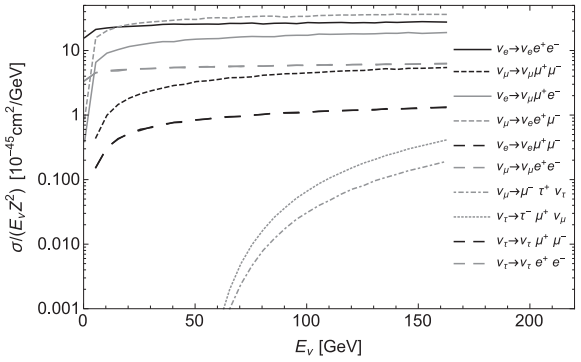

FIG. 3. σ/E_ν trident cross sections normalized by Z^2 for various SM flavors as a function of the incoming neutrino energy on a lead target (SHiP).

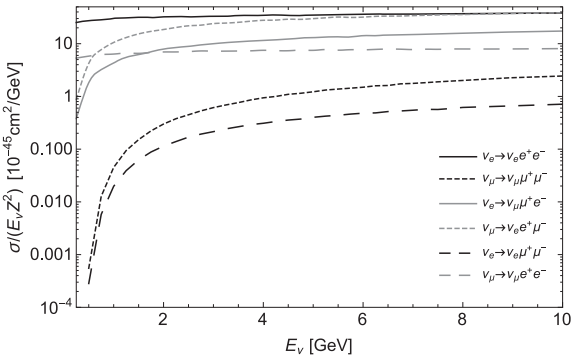

FIG. 4. σ/E_ν trident cross sections normalized by Z^2 for various SM flavors as a function of the incoming neutrino energy on an argon target (DUNE).

IV. DISCUSSION AND ANALYSIS

The general features of Sec. III can be understood qualitatively by considering Eq. (4) and Table I. First, we note that every cross section is proportional to the combination $|C_V|^2 + |C_A|^2$ appearing in Eq. (4). In the SM, this is maximal in the case of W mediated interactions, intermediate for W + Z mediated interactions, and minimal for Z mediated interactions. The W exclusive channel corresponds to scattering events where the incoming and outgoing neutrinos belong to different lepton generations, and thus these channels will be more probable. Another dominant feature controlling the relative size of cross sections is related to the masses of the outgoing leptons. This dictates the size of the logarithmic enhancement coming from the low-Q^2 phase space. This is a feature of the IR divergence arising from the photon propagator, which is regulated by the finite masses of the charged leptons. Finally, the rates quoted in Tables II to IV are further influenced by beam luminosity and so tend to favor incident muon configurations, except at the DUNE far detector, where they favor incident electron neutrinos.

These qualitative features suggest that $\nu_\mu \to \nu_e \mu^- e^+$ would serve as the dominant production mode at both the DUNE near detector and SHiP. Examining Tables II and III, this is indeed the case. It is a CC-exclusive process (high axial-vector couplings), and it benefits from the large flux of muon neutrinos and from the logarithmic enhancement afforded by the low electron mass. This final statement is most important at DUNE due to its lower $\langle E_\nu \rangle$, which makes it sensitive to muon-mass threshold effects. For diffractive processes, the sensitivity of the cross section to the charged-lepton masses is weakened due to the lower bound $Q_{\min}$ in Eq. (6). This accounts for the difference in ordering of rates between the coherent and diffractive contributions to the cross section found in Tables II to IV. At DUNE, this results in an enhancement of the cross section by a factor of 35 when compared to the production mode $\nu_\mu \to \nu_\mu \mu^+ \mu^-$, which was observed at CHARM-II, CCFR, and NuTeV [1,3,17]. No dedicated search was carried out for electron production in trident modes at these experiments. The detector technology typically consisted of interwoven layers of heavy element materials to induce neutrino interactions, followed by calorimeters to measure the final lepton states. Electrons create showers and scatter much more in these layers, as opposed to muons which tend to follow a straight trajectory until the muon spectrometer. It was thus much more difficult to impose vertex requirements on electrons, which is an integral part of the trident analysis. Neutrino detector technology has greatly evolved since then, and it is now feasible to consider mixed flavor trident channels.

The lifetime expected event counts for $\mu^+\tau^-$ and $\mu^-\tau^+$ production are both approximately unity. Given the uncertain run time and technical specifications of SHiP. it is possible that tridents containing tau leptons will occur; however, the rates are sufficiently low that it is not clear at what level of statistical significance these can be observed, especially after applying necessary cuts. Our analysis suggests that these events are most likely to occur for intermediate momentum transfers (i.e. in the diffractive regime). Our deep-inelastic analysis revealed high-Q^2 trident production to be extremely suppressed for all flavors, including tau leptons. ν_τ induced electon-muon pairs may be observable; however, due to the much higher flux of ν_μ's, this channel will be dominated by ν_μ induced events with identical charged-lepton final states, which will leave an indistinguishable signature in the detector.

In the case of the DUNE Collaboration, the size of the near detector is currently being planned such that it can obtain approximately ten times the statistics of the far detector, allowing for a reduction in the systematic uncertainties of the neutrino beam. Our results show that, even for near detector masses that minimally satisfy this requirement, trident production should be detectable. Given the large beam intensity at the near detector, every additional unit of detector mass represents a fantastic return on investment from the perspective of rare neutrino processes such as trident production. Pushing from hundreds to thousands of events would lower the statistical error to the level of a few percent and could potentially allow for trident production to act as a complimentary beam characterization tool. This is alluring because trident production is only sensitive to the target nucleus's electric form factor, in contrast to CC events where uncertainties in the axial form factor still introduce significant systematic effects.

While interesting in its own right as a test of the Standard Model, neutrino trident production can also act as a significant background in the search for new physics. This is because of its qualitative similarities to processes involving lepton flavor violation, which is a signature of many BSM models. Our estimated rates also suggest that both SHiP and the DUNE near detector can be used to constrain BSM physics; comparison with the number of events identified by the CCFR and CHARM-II Collaboration in the dimuon channel alone demonstrates that both SHiP and DUNE are competitive with these previous experiments. With access to flavor dependent final states, however, we believe these experiments can do much better. For example, the Z' coupling to $L_\mu - L_\tau$ considered in Ref. [4] influences both $\nu_\mu \to \nu_\mu \mu^+ \mu^-$ and $\nu_\mu \to \nu_\mu e^+ e^-$. Due to the minimal size of $|C_V|^2 + |C_A|^2$ for e^+e^- production (due to Z-exclusive mediation), this process will experience an even greater relative sensitivity to new physics, albeit in a first-generation lepton channel.

Although the qualitative features discussed earlier are sufficient to understand the most prominent aspects of our analysis, a closer examination of Figs. 3 and 4 reveals another feature, which is initially surprising. The rates for processes which seem to be related by an exchange of flavor indices have different cross sections. This effect is

$\mathcal{O}(1)$ and independent of energy (see Fig. 3, $\nu_\mu \to \nu_\tau\tau^+\mu^-$ vs $\nu_\tau \to \nu_\mu\mu^+\tau^-$ for example). This would seem to suggest a violation of lepton universality; however, a closer examination reveals that the chiral structure of the outgoing leptons is not equivalent, with the amplitudes for production into inequivalent configurations being proportional to the square of the heaviest lepton mass. Still, this effect is surprising given that it is independent of energy, and naively one would expect that at sufficiently high center-of-mass energies the effect would be suppressed by m_ℓ^2/S with S the Mandelstam variable for the neutrino-nucleus interaction. This is not the case for trident production because the cross section is dominated by the low-Q^2 region of phase space. To understand this, we turn to the EPA and more specifically Eq. (3). We see that the integral over s has an IR cutoff of $m_{ij}^2 = (m_{\ell^+} + m_{\ell^-})^2$, and so in this regime, we find an $\mathcal{O}(m_{ij}^2/s) \sim \mathcal{O}(1)$ contribution to the cross section, which will be present even for arbitrarily high E_ν.

To understand why the chiral structure of the amplitude has a significant impact on the amplitude, we must consider both the infrared divergence of the photon mediator and the constraints imposed by conservation of angular momentum. Consider the center-of-mass frame for the photon-neutrino pair. To saturate the lower bound of the integral over s in Eq. (3), we must produce the lepton pair at rest and have the neutrino redshift to an arbitrarily small energy $E_\nu' = \epsilon$; this also forces the lepton pair to carry equal and opposite momentum. It is, however, difficult to understand the implications of chirality in this frame, because in this frame the lepton pair is nonrelativistic and we cannot freely interchange helicity and chirality.

To solve this problem, we can appeal to Lorentz invariance and perform the same analysis in a boosted frame in which the lepton pair is highly relativistic. The boost should be directed in the direction of the lepton pair's infinitesimal three-momentum. This boost will further redshift the outgoing neutrino, but it will not change its direction. We would like to check if this configuration conserves angular momentum, and the answer to this question is dependent on the initial polarization of the incident photon (the neutrino's polarization is fixed because of its definite chirality), which in turn determines the initial angular momentum. The two possibilities are $S = 1/2$ and $S = 3/2$.

As shown in Fig. 5, the outgoing states for the two configurations have different chirality [a triplet of left-handed leptons (LLL) vs a triplet of right-handed leptons (LRR)]. In the case of $S = 1/2$, this has no effect on the configuration discussions above; however, in the case of $S = 3/2$, where the spin of the neutrino and photon are aligned, the LLL configuration is forbidden, while the LRR configuration is allowed. This is because in our boosted frame, where chirality is equivalent to helicity, in order to obtain $S = 3/2$ for the LLL configuration, all three particles would have to travel in the same direction, which would violate the conservation of momentum. Thus, only the LRR, and not the LLL, configuration satisfies all the necessary conservation laws in the low-Q^2 region of phase space that dominates Eq. (3).

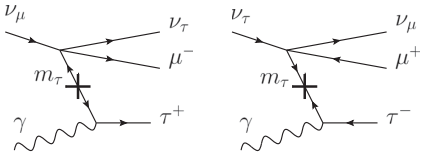

FIG. 5. Inequivalent contributions to the processes $\nu_\mu \to \nu_\tau\tau^+\mu^-$ (left) and $\nu_\tau \to \nu_\mu\mu^+\tau^-$ (right) in the limit of $m_\mu \to 0$. Note that the chiral structure of the weak interaction results in a triplet of left-handed leptons for incident ν_μ and a right-handed lepton pair with a left-handed neutrino for incident ν_τ. The fermions are two component spinors of definite chirality. Diagramatic conventions are from Ref. [20] with arrows denoting chirality.

V. CONCLUSIONS AND OUTLOOK

We have demonstrated that so-far-unobserved neutrino trident processes are within reach with the planned DUNE and SHiP experimental collaborations. The DUNE Collaboration may be able to enhance production modes, some of which we currently estimate to only yield 1–10 events in the experiment's lifetime (e.g. $\nu_\mu \to \nu_\mu\mu^+\mu^-$), by increasing the mass of the relatively small near detector. Even with the current proposed designs, both collaborations are maximally sensitive to the modes $\nu_\mu \to \nu_e\mu^-e^+$ and $\bar{\nu}_\mu \to \bar{\nu}_e\mu^+e^-$. We believe that backgrounds for these searches will be low, especially given the vertex resolution of both experiments [6,7].

In addition to our direct application to the DUNE and SHiP collaborations, we also present $\sigma(E_\nu)$ for the coherent scattering regime, allowing for future analyses with more precise luminosity estimates. We present a similar plot in Fig. 7 in case high momentum-transfer trident production is of future interest. We have considered all possible combinations of lepton flavor final states and have presented only processes with nonzero lifetime event counts. This work is complementary to that found in Ref. [11], in which differential distributions with respect to the lepton pair's invariant mass are plotted in the coherent regime. Additionally, we have demonstrated a method for treating neutrino trident production on the parton level, which requires some slight modifications to the standard treatment. This revealed high-Q^2 trident production is untenable as one would naively expect.

Neutrino trident production is a proven tool in the testing of the SM and constraining BSM physics, and with improved detector designs, it is important to harness the full capabilities of next-generation neutrino experiments. Our analysis suggests that both SHiP and DUNE will be able to observe trident production. We believe that with

these experiments on the horizon the future is bright for studying trident production and other rare neutrino processes and that the study of these processes should be incorporated into the physics programs of both experiments.

ACKNOWLEDGMENTS

We are very grateful to Maxim Pospelov for suggesting mixed flavor trident production and its applicability to future intensity frontier experiments. Additionally, we would like to thank him for his continued guidance throughout this research. We would also like to thank Itay Yavin for his help in the treatment of phase space. Finally, we thank Chien-Yi Chen and Richard Hill for useful discussions. This research was supported in part by Perimeter Institute for Theoretical Physics. Research at Perimeter Institute is supported by the Government of Canada through the Department of Innovation, Science and Economic Development and by the Province of Ontario through the Ministry of Research and Innovation. This research was also supported by funds from the National Science and Engineering Research Council of Canada, the Ontario Graduate Scholarship program, and the Early Research Awards program of Ontario.

APPENDIX A: THREE-BODY PHASE SPACE (EPA)

For the purposes of the EPA, the phase space integrals are performed over the three-body phase space of the leptons. Ultimately, this three-body phase space is embedded in the full four-body one, and so we will use the results of this section in the proceeding one. We denote the center-of-mass energy for the photon-neutrino collision by s; additionally, we define the quantities $P = p_+ + p_-$ and $\ell = P^2$. We begin by decomposing the three-body phase space using the identity below:

$$\mathrm{d}\Phi_3(p_+, p_-, k_2) = \frac{\mathrm{d}\ell}{2\pi}\Phi_2(k_2, P)\Phi_2(p_+, p_-). \tag{A1}$$

Each two-body phase space can be expressed as

$$\mathrm{d}\Phi_2(q_1, q_2) = \bar{\beta}(q_1, q_2)\frac{\mathrm{d}\Omega}{32\pi^2} \tag{A2}$$

with the definition

$$\bar{\beta}(q_1, q_2) = \sqrt{1 - \frac{2(q_1^2 + q_2^2)}{(q_1 + q_2)^2} + \frac{(q_1^2 - q_2^2)^2}{(q_1 + q_2)^4}}. \tag{A3}$$

An important case is when $q_1^2 = 0$. In this scenario, the factor simplifies to $\bar{\beta} = 1 - \frac{q_2^2}{(q_1+q_2)^2}$. In our decomposition above, $\bar{\beta}(k_2, P) = 1 - \ell/s$. First, we choose to evaluate $\mathrm{d}\Phi_2(P, k_2)$ in the center-of-mass frame of the reaction. This allows us to parametrize the phase space as written in Eq. (A2). We can perform the azimuthal integration by appealing to symmetry, and we are left only with $\mathrm{d}\cos\theta_{CM}$. This can conveniently be expressed in terms of the Lorentz-invariant t defined via

$$t = 2q_\mu(k_1 - k_2)^\mu = \frac{1}{2}(s + \ell + (\ell - s)\cos\theta_{CM}). \tag{A4}$$

This definition leads to the differential relationship $\mathrm{d}t = \frac{1}{2}(\ell - s)\mathrm{d}\cos\theta$. Thus, we can simplify our three-body phase space integral by applying the identity $\bar{\beta}(k_2, P)d\cos\theta = -\frac{2}{s}\mathrm{d}t$. This leaves us with the second phase space integral. This is most easily evaluated in the frame where P_μ has vanishing three-momentum. In this frame, there is no guarantee of azimuthal symmetry in the matrix element, and so we must integrate over both polar angles. We are left with the expression quoted in Ref. [4],

$$\mathrm{d}\Phi_3(k_2, p_+, p_-) = \frac{1}{2}\frac{1}{(4\pi)^2}\frac{\mathrm{d}\ell}{2\pi}\bar{\beta}(p_+, p_-)\frac{\mathrm{d}t}{2s}\frac{\mathrm{d}\Omega}{4\pi}, \tag{A5}$$

where we denote the angular integral over the muon pair, performed in the frame where $P = (\sqrt{\ell}, 0, 0, 0)$ by $\mathrm{d}\Omega$. The limits of integration for t are given by $\ell < t < s$. This gives the expression for the photon-neutrino cross section as

$$\sigma_{\gamma\nu} = \frac{1}{2s}\frac{1}{2}\frac{1}{(4\pi)^2}\int_{m_{jk}^2}^{s}\frac{\mathrm{d}\ell}{2\pi}\bar{\beta}_\pm(\ell)\int_\ell^s\frac{\mathrm{d}t}{2s}\int\frac{\mathrm{d}\Omega}{4\pi}|\overline{\mathcal{M}}|^2_{\gamma\nu}, \tag{A6}$$

where $m_{jk} = m_j + m_k$, and $|\overline{\mathcal{M}}|^2 = 1/2\sum_{\mathrm{pol}}|\mathcal{M}|^2$.

To obtain the full cross section, this must be weighted against the probability for creating a photon in the Coulomb field of a nucleus, given in Refs. [4,5]. This leads to

$$\sigma_{N\nu} = \frac{Z^2\alpha}{\pi}\int_{m_{jk}^2}^{s}\frac{\mathrm{d}s}{s}\sigma_{\gamma\nu}(s)\int_{(s/2E_\nu)^2}^{\infty}\frac{\mathrm{d}Q^2}{Q^2}F^2(Q^2), \tag{A7}$$

where $\sqrt{S}$ denotes the neutrino-nucleus center-of-mass energy. In practice, the form factor of the nucleus $F(Q^2)$ cuts this integral off near $s_{\max} \approx 2E_\nu\Lambda_{\mathrm{QCD}}/A^{1/3}$. In our calculations for the coherent regime (Sec. II B 1), we used the Woods-Saxon form factor

$$F_{\mathrm{WS}}(Q^2) = \frac{1}{N}\mathcal{F}\left\{\frac{V_0}{1 + \exp(\frac{r - r_0A^{1/3}}{a})}\right\} \tag{A8}$$

with $\mathcal{F}$ denoting the Fourier transform with respect to r and N denoting a normalization factor given by the volume integral over the nuclear charge distribution [21]. The various parameters are set as $r_0 \approx 1.126$ fm, $a \approx 0.523$ fm, and $V_0 = (4\pi A r_0^3/3)^{-1}$. Different choices of the form factor modify the result on the 10% level.

For the diffractive regime, we used the electric dipole fit for the proton's Dirac form factor found in Refs. [9,13,22]. Due to the quasielastic nature of the scattering, the Pauli form factor's contribution is suppressed. The explicit expression is given by

$$F_{\rm dip}(Q^2) = \frac{G_{\rm dip}(Q^2) + \tau\xi G_{\rm dip}(Q^2)}{1+\tau}, \tag{A9}$$

where $\tau = Q^2/4M^2$ with $M = (m_p + m_n)/2$ and $\xi = (\mu_p - \mu_n)/\mu_N \approx 4.7$ is the difference in magnetic moments between the proton and the neutron measured in units of the nuclear magneton. The dipole fit is given by

$$G_{\rm dip}(Q^2) = \left(1 + \frac{Q^2}{0.71\ \text{GeV}^2}\right)^{-2}. \tag{A10}$$

APPENDIX B: FOUR-BODY PHASE SPACE (DIS)

1. Parton-neutrino collision

We now consider the decomposition of the four-body phase space. This will involve a reduction to the previously analyzed three-body case; however, there will be some difference thereafter because of the loss of azimuthal symmetry in $\Phi_2(P, k_2)$.

We begin by emphasizing a change in notation. The center-of-mass energy for the parton-neutrino collision is denoted S, we introduce the four-vector $R = k_2 + p_+ + p_-$ and its invariant mass $L = R^2$, and we maintain the previous definition of $P = p_- + p_+$. We can now decompose the four-body phase space as shown schematically in Fig. 6 and more precisely below:

$$\begin{aligned} &d\Phi_4(p_+, p_-, h_2, k_2) \\ &\quad = \frac{dL}{2\pi} d\Phi_2(R, h_2) d\Phi_3(p_+, p_-, k_2) \\ &\quad = \frac{d\ell}{2\pi}\frac{dL}{2\pi} d\Phi_2(R, h_2)\Phi_2(k_2, P)\Phi_2(p_+, p_-). \end{aligned} \tag{B1}$$

The first two-body phase space $\Phi_2(h_2, R)$ inherits the azimuthal symmetry of the parton-neutrino collision, and in direct analogy with Eq. (A4), we introduce the variable T defined via

$$T = 2h_1^\mu (k_1 - h_2)_\mu = \frac{1}{2}[S + L + (L - S)\cos\theta_h]. \tag{B2}$$

The final pair of two-body phase spaces do not inherit the azimuthal symmetry, and so we do not attempt to further simplify them. We therefore evaluate $d\Phi_2(k_2, P)$ and $d\Phi_2(p_+, p_-)$ in their respective rest frames. The angles of the charged-lepton frame are labelled θ and ϕ, while those of $\Phi_2(k_2, P)$ are labelled θ' and ϕ'. With these variables, the four-body phase space can be written

$$d\Phi_4 = \frac{dL}{2\pi}\frac{4\pi}{32\pi^2}\frac{dT}{S}\frac{d\ell}{2\pi}\bar{\beta}(k_2, P)\bar{\beta}(p_+, p_-)\frac{d\Omega'}{32\pi^2}\frac{d\Omega}{32\pi^2}. \tag{B3}$$

Keeping in mind that the Lorentz-invariant flux factor $\mathcal{F}$ for massless initial states is given by $\mathcal{F} = 2S$, we can express the parton cross section as

$$\sigma_{h\nu}(S) = \frac{1}{2S}\int_{m_{jk}^2}^{S}\frac{dL}{2\pi}\int_L^S \frac{2\pi}{32\pi^2}\frac{2dT}{S}\int_{m_{jk}^2}^{L}\frac{d\ell}{2\pi}\bar{\beta}(k_2, P)\bar{\beta}(p_+, p_-)\int\frac{d\Omega}{32\pi^2}\int\frac{d\Omega'}{32\pi^2}|\overline{\mathcal{M}}|^2_{h\nu}. \tag{B4}$$

2. Hadron neutrino cross section

We now connect our partonic cross section to the hadronic cross section via the formalism of deep-inelastic scattering. We introduce the new variable S_H defined by

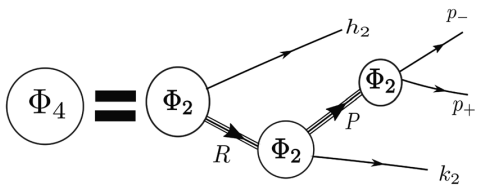

FIG. 6. Schematic depiction of the four-body phase space decomposition into three two-body phase spaces. Note the three-body phase space decomposition for the EPA is obtained by considering only the final two phase spaces in the diagram.

$\xi S_H = S$ and is given by $S_H = 2E_\nu M_N$ in the lab frame. Unlike in textbook treatments of deep-inelastic scattering, we cannot integrate ξ over the full interval [0, 1] because we require a minimum amount of energy to produce the pair of charged leptons (i.e. $\xi \approx 0$ is kinematically forbidden). Additionally, we would like to ensure that we do not double-count amplitudes already included in the EPA, and so we include a cut on the minimum amount of four-momenta transfer to the nucleus $Q > Q_{\rm min}$.

To impose this cut, it is easiest to change from the variable T to the variable $U = Q^2 = |q^2|$. If we place a cut on the momentum transfer $U > Q^2_{\rm min}$, then this changes the bounds of integration in Eq. (B4). We chose $Q_{\rm min} = 1$ GeV to ensure we are not double-counting amplitudes. However, with this scheme, we include a parametric regime in which hadronic resonances can be very

important. Although the description in terms of partons may capture some of the essential features of hadron production, it is probable that the DIS formalism underestimates the rates, because it does not incorporate resonance conditions.

The effects of a cut on momentum transfer can be seen by noting that $U = S - T$ and that the bounds of integration require $S > T > L$. The smallest L, and by proxy T, can be is m_{jk}^2, which implies that $U = S - T < S - m_{jk}^2$. Combining this with the condition that $U > Q_{\min}$ leads to $S - m_{jk}^2 > Q_{\min}^2$. Finally, this can be converted into a minimum bound on ξ given by

$$\xi \geq \frac{Q_{\min}^2 + m_{jk}^2}{S_H}. \tag{B5}$$

Finally, we note that, depending on the magnitude and direction of the individual leptons, Q^2 could range from being very small to $S - m_{jk}^2$, and so we must include the parton distribution functions inside the integral over U. This leads to our final expression for the nucleon-neutrino cross section,

$$\sigma_{H\nu} = \sum_h \int_{\xi_{\min}}^1 \frac{\mathrm{d}\xi}{2\xi S_H} \int_{m_{jk}^2}^{L_{\max}} \frac{\mathrm{d}L}{2\pi} \int_{Q_{\min}^2}^{\xi S_H - L} \frac{1}{8\pi} \frac{\mathrm{d}U}{\xi S_H} \int_{m_{jk}^2}^{L} \frac{\mathrm{d}\ell}{2\pi} \int \frac{\mathrm{d}\Omega}{32\pi^2} \frac{\mathrm{d}\Omega'}{32\pi^2} |\overline{\mathcal{M}}|_{h\nu}^2 \bar{\beta}(k_2, P) \bar{\beta}(p_+, p_-) f_h^{(H)}(\xi, U), \tag{B6}$$

where h runs over all the partons in the given nucleon $H \in \{n, p\}$ (either neutrons or protons), $L_{\max} = \xi S_H - Q_{\min}^2$, $\xi_{\min}$ saturates the bound in Eq. (B5), and $f_h^{(H)}(\xi, Q)$ is the parton distribution function for the parton h in H. To obtain the neutrino-nucleus cross section, a simple weighted sum of individual nucleon cross sections was used:

$$\sigma_{A\nu} = Z\sigma_{p\nu} + (A - Z)\sigma_{n\nu}. \tag{B7}$$

APPENDIX C: LUMINOSITY ESTIMATES

1. SHiP

For the purposes of calculating expected rates at SHiP, we relied on Ref. [6]; specifically Fig. 5.25 and Table 2.3. These quote the number of expected charged-current events in the detector. To convert this into a neutrino luminosity, we simply divided by the charged-current cross section which we took to be given by

$$\sigma_{CC} = A\left(\frac{E_\nu}{\mathrm{GeV}}\right) \begin{cases} 6.75 \times 10^{-39}\ \mathrm{cm}^{-2}(\nu) \\ 3.38 \times 10^{-39}\ \mathrm{cm}^{-2}(\bar{\nu}) \end{cases} \tag{C1}$$

with the braced numbers referring to incident neutrinos and antineutrinos respectively. To determine the experiment's lifetime integrated luminosity, we used the number of CC events from Table 2.3 of Ref. [6], while the energy spectrum was taken from Fig. 5.25. Finally, we multiplied by the detector's efficiency (which we took to be 90% for each of the final state leptons), leading to Eq. (9).

2. DUNE

The DUNE Collaboration's far and near detectors are treated separately in their proposals, with a heavier emphasis on the far detector. As a result, there is no published neutrino spectrum for the near detector; however, both detectors have lifetime expected event counts. We therefore had to infer the near detector spectrum from that of the far detector and then normalize our results to reproduce the lifetime rates quoted in Table 6.1 of Ref. [7].

To link the beam luminosity in the far detector to those in the near detector, we also adjusted the various flavors' luminosities to account for oscillation effects. All ν_e appearances at the far detector were assumed to stem from ν_μ at the near detector, while $\nu_e + \bar{\nu}_e$ background in the far detector was assumed to represent the full flux of first-generation neutrinos at the near detector up to geometric losses due to beam spread.

Additionally, the CC rates in the DUNE proposal at the near detector are quoted per 10^{20} POT and 1 ton of detector mass. The far detector rates are quoted assuming 150 kt-MW-yr. This number assumes a 40 kt far detector and that 1.2 MW of beam power corresponds to 1.1×10^{21} POT/yr. We therefore multiply the event counts in Table 6.1 in Ref. [7] by

$$\frac{1.1 \times 10^{21}\ \mathrm{POT/yr}}{1.2\ \mathrm{MW}} \times 850\ \text{kt-MW-yr} \times \frac{0.1\ \text{ton}}{40\ \text{kt}}, \tag{C2}$$

where 850kt-MW-yr is the exposure at the far detector in the lifetime of DUNE given the optimized design and 0.1 ton is the mass of the near detector.

Next, we consider the details of the far detector. For this, we use Figs. 3.5 and 3.29 and Table 3.5. Table 3.5 and Fig. 3.5 are in correspondence with one another, and we quote their results for an exposure of 150kt-MW-yr. They specify different rates for the running of the experiment in the neutrino and antineutrino modes; we presume each mode constitutes half of the experiment's lifetime. We therefore adjust the rates quoted in Table 3.5 and Fig. 3.5 of Ref. [7] by a factor of

$$\frac{850\ \text{kt-MW-yr}}{150\ \text{kt-MW-yr}} \times \frac{1}{2} \tag{C3}$$

to obtain the lifetime event rate for the far detector. The spectrum is given in Fig. 3.29 and is quoted here in units of CC events/GeV/kT/yr. The experiment is set to obtain an exposure of 300 kt-MW-yr at 1.07 MW and then 550 kt-MW-yr at 2.14 MW. Additionally, the energy bin width of the plot is 0.25 GeV, and so we multiply the spectrum of Fig. 3.29 of Ref. [7] by a factor of

$$\frac{0.25\ \text{GeV}}{1\ \text{bin}} \left(\frac{300\ \text{kt-MW-yr}}{1.07\ \text{MW}} + \frac{550\ \text{kt-MW-yr}}{2.14\ \text{MW}} \right). \tag{C4}$$

Finally, in Fig. 3.29, the individual CC-event rates of ν_e and $\bar{\nu}_e$ are not given, but their sum is given. We assumed the relative ratio of neutrinos to antineutrinos was equal to the appearance rates quoted in Table 3.5 of Ref. [7]. Although the background neutrino rates are much smaller than the oscillation signal, they provide the dominant contribution at the near detector. The production fractions of K^+ and K^- kaons, denoted $R_{K^\pm}$, have to be compared with those of π^+ and π^-, given as $R_{\pi^\pm}$.

We therefore assume that at the far detector the relative components of the $\nu_e + \bar{\nu}_e$ background are given by

$$N_{\nu_e}^{(\text{Bkg})} = \frac{R_{\pi^+}}{R_{K^+}} \frac{N_{\nu_e}^{(\text{osc})}}{N_{\text{tot}}^{(\text{osc})}} N_{\text{tot}}^{(\text{Bkg})} \tag{C5a}$$

$$N_{\bar{\nu}_e}^{(\text{Bkg})} = \frac{R_{\pi^-}}{R_{K^-}} \frac{N_{\bar{\nu}_e}^{(\text{osc})}}{N_{\text{tot}}^{(\text{osc})}} N_{\text{tot}}^{(\text{Bkg})}. \tag{C5b}$$

We then assume the first-generation component at the near detector is the progenitor of the full background signal at the far detector. Equivalently, we estimate the number of electron and antielectron events at the near detector to be proportional to N_{Bkg} at the far detector with an overall normalization that is consistent with geometric losses. To find the geometric loss factor, we compared the rates for ν_μ CC events quoted in Table 6.1 of Ref. [7] with the CC events from the ν_μ background signal and ν_e appearance signal quoted in Table 3.5 and Fig. 3.5 of Ref. [7]. Our beam spectrum at the far detector was taken from Figs. 3.29 and 3.5 of Ref. [7].

APPENDIX D: DEEP-INELASTIC SCATTERING RESULTS

See Figs. 7 and 8.

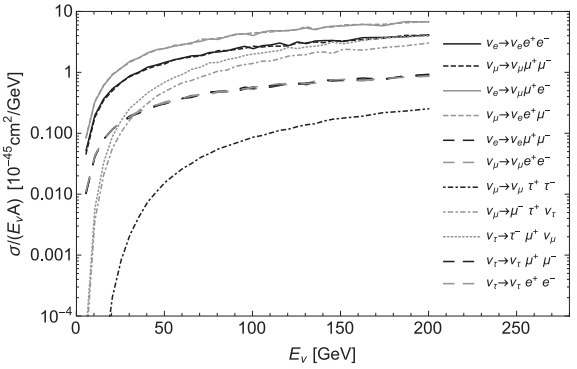

FIG. 7. σ/E_ν trident DIS cross sections per nucleon for various SM flavors as a function of the incoming neutrino energy on a lead target (SHiP).

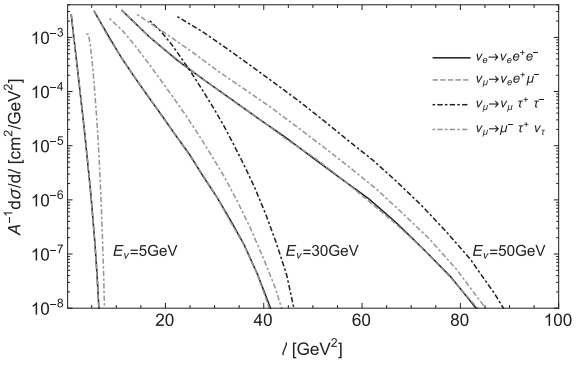

FIG. 8. Normalized $d\sigma/dl$ for a variety of DIS processes at SHiP, where $l = (p_{l^+} + p_{l^-})^2$. Energies are in GeV.

[1] D. Geiregat and G. Wilquet (CHARM II Collaboration), First observation of neutrino trident production, Phys. Lett. B **245**, 271 (1990).

[2] S. R. Mishra *et al.*, Neutrino tridents and W-Z interference, Phys. Rev. Lett. **66**, 3117 (1991).

[3] T. Adams (NuTeV Collaboration), Neutrino Trident Production from NuTeV, arXiv:hep-th/9811012.

[4] W. Altmannshofer, S. Gori, M. Pospelov, and I. Yavin, Neutrino Trident Production: A Powerful Probe of New Physics with Neutrino Beams, Phys. Rev. Lett. **113**, 091801 (2014).

[5] R. Belusevic and J. Smith, W-Z interference in ν-nucleus scattering, Phys. Rev. D **37**, 2419 (1988).

[6] M. Anelli *et al.* (SHiP Collaboration), A facility to Search for Hidden Particles (SHiP) at the CERN SPS, arXiv:1504.04956.

[7] R. Acciarri *et al.*, Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE) Conceptual Design Report Volume 2: The Physics Program for DUNE at LBNF, arXiv:1512.06148.

[8] I. Yavin (private communication).

[9] J. A. Formaggio and G. P. Zeller, From eV to EeV: Neutrino cross sections across energy scales, Rev. Mod. Phys. **84**, 1307 (2012).

[10] A. Drukier and L. Stodolsky, Principles and applications of a neutral-current detector for neutrino physics and astronomy, Phys. Rev. D **30**, 2295 (1984).

[11] J. Løvseth and M. Radomski, Kinematical distributions of neutrino-produced lepton triplets, Phys. Rev. D **3**, 2686 (1971).

[12] M. I. Vysotsky, I. V. Gaidaenko, and V. A. Novikov, On lepton-pair production in neutrino-nucleus collisions, Phys. Lett. B **65**, 1676 (2002).

[13] R. W. Brown, R. H. Hobbs, J. Smith, and N. Stanko, Intermediate boson. III. Virtual-boson effects in neutrino trident production, Phys. Rev. D **6**, 3273 (1972).

[14] A. D. Martin, W. J. Stirling, R. S. Thorne, and G. Watt, Parton distributions for the LHC, Eur. Phys. J. C **63**, 189 (2009).

[15] J. A. M. Vermaseren, New features of FORM, arXiv:math-ph/0010025.

[16] K. A. Olive *et al.* (Particle Data Group), Review of particle physics, Chin. Phys. C **38**, 090001 (2014).

[17] D. Geiregat *et al.*, Calibration and performance of the CHARM-II detector, Nucl. Instrum. Methods Phys. Res., Sect. A **325**, 92 (1993).

[18] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H.-S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations, J. High Energy Phys. 07 (2014) 079.

[19] W. Bonivento *et al.*, Proposal to Search for Heavy Neutral Leptons at the SPS, arXiv:1310.1762.

[20] H. K. Dreiner, H. E. Haber, and S. P. Martin, Two-component spinor techniques and Feynman rules for quantum field theory and supersymmetry, Phys. Rep. **494**, 1 (2010).

[21] U. D. Jentschura and V. G. Serbo, Nuclear form factor, validity of the equivalent photon approximation and Coulomb corrections to muon pair production in photon-nucleus and nucleus-nucleus collisions, Eur. Phys. J. C **64**, 309 (2009).

[22] C. F. Perdrisat, V. Punjabi, and M. Vanderhaeghen, Nucleon electromagnetic form factors, Prog. Part. Nucl. Phys. **59**, 694 (2007).

Chapter 6

Probing new charged scalars with neutrino trident production

6.1 Preface

In the previous chapter, we've calculated rates for mixed-flavor trident production in the context of the SM, and shown that certain of these modes may indeed be very large compared to the conventional $\mu^+\mu^-$ mode. In this section, we investigate if these topologies, never previously observed, can do a good job at constraining models of new physics. This investigation is partially motivated by the successful application of the $\mu^+\mu^-$ channel in constraining models of Z' gauged under $L_\mu - L_\tau$ [43]. A natural avenue to investigate are models involving leptophilic charged scalars since they can show non-diagonal flavor structure to leptons. Due to their universal interactions with photons however, charged scalars are constrained via LEP data to have masses above $80 - 100$ GeV, regardless of the details of their interactions with neutrinos and charged leptons. In this work, we perform a systematic overview of the bounds one can achieve in these various mixed-flavor channels, assuming only the irreducible SM mixed-flavor NTP backgrounds. We compare these projected limits to bounds in the literature coming from universality and flavor violating processes. We find that under selected assumptions for the nature of the coupling constants in the new

physics model, NTP can be competitive to existing constraints. If however we consider implementations of these charged scalars in the literature that explain neutrino masses, NTP becomes uncompetitive.

 It has been published in PRD as

- Gabriel Magill and Ryan Plestid, "Probing new charged scalars with neutrino trident production," Phys.Rev. **D97** (2018), arXiv:1710.08431 [hep-ph].

Contribution of author (reproduced from Declaration of Authorship): Both authors contributed significantly to the calculations, writing and editing of the paper.

6.2 Paper

PHYSICAL REVIEW D **97,** 055003 (2018)

Probing new charged scalars with neutrino trident production

Gabriel Magill[*] and Ryan Plestid[†]

Department of Physics and Astronomy, McMaster University, Hamilton, Ontario, Canada and Perimeter Institute for Theoretical Physics, Waterloo, Ontario, Canada

(Received 20 November 2017; published 5 March 2018)

We investigate the possibility of using neutrino trident production to probe leptophilic charged scalars at future high intensity neutrino experiments. We show that under specific assumptions, this production process can provide competitive sensitivity for generic charged scalars as compared to common existing bounds. We also investigate how the recently proposed mixed-flavor production—where the two oppositely charged leptons in the final state need not be muon flavored—can give a 20%–50% increase in sensitivity for certain configurations of new physics couplings as compared to traditional trident modes. We then categorize all renormalizable leptophilic scalar extensions based on their representation under $\mathrm{SU}(2) \times \mathrm{U}(1)$, and discuss the Higgs triplet and Zee-Babu models as explicit UV realizations. We find that the inclusion of additional doubly charged scalars and the need to reproduce neutrino masses make trident production uncompetitive with current bounds for these specific UV completions. Our work represents the first application of neutrino trident production to study charged scalars. Additionally, it is the first application of mixed-flavor trident production to study physics beyond the standard model more generally.

DOI: 10.1103/PhysRevD.97.055003

I. INTRODUCTION AND MOTIVATION

Neutrino oscillation experiments provide conclusive evidence that the standard model (SM) is incomplete. Many unresolved anomalies—the proton radius puzzle [1,2], the anomalous magnetic moment of the muon [3,4], and the Liquid Scintillator Neutrino Detector anomaly [5]—can be interpreted as providing hints into beyond the SM (BSM) physics, especially for heavy leptons where constraints are typically weaker. Scalar extensions of the SM have been proposed as solutions to all of these anomalous measurements [6–9]. Currently, most constraints on the scalar sector come from low energy observables and high energy colliders [10,11]. In contrast, high intensity mid energy neutrino experiments have remained relatively uninvestigated. Consequently, new tools sensitive to interactions between scalars and neutrinos/heavy leptons provide a complimentary probe of BSM physics.

Neutrino trident production (NTP) represents a natural candidate for studying couplings to an extended scalar sector given the successful application of NTP to models with an Abelian Z' coupled to $L_\mu - L_\tau$ [12]. Using data from the beam dump experiments CHARM-II and CCFR [13,14] the authors of Ref. [12] were able to probe previously unexplored parameter space, including part of the favored region for the resolution of the $(g-2)_\mu$ anomaly. As demonstrated in Ref. [15], the upcoming beam dump experiments SHiP and DUNE [16,17] are sensitive to many previously unmeasured neutrino trident channels which contain mixed-flavor leptons in the final state. With these exciting new prospects the possibility of NTP serving as a powerful probe of scalar extensions seems highly probable. Furthermore, given the mounting interest in precision neutrino physics, NTP may find applications at other future neutrino experiments, in particular, Fermilab's short-baseline neutrino program [18].

NTP involves the creation of a lepton pair via a high energy neutrino scattering coherently (diffractively) with a nucleus (nucleon) as shown in Fig. 1. This production mechanism is subdominant to charged-current (CC) scattering, in large part due to the extra α^2 fine-structure suppression in its cross section; for 50 GeV neutrinos scattering coherently on lead producing a $\mu^+\mu^-$ final state, we expect one trident event for every 10^5 CC events [19]. As discussed in Ref. [15], this scaling depends largely on the flavors of the final state lepton pair, with event rates being 40 times larger in the case of $e^+\mu^-$ production at DUNE. This is due to the absence of W-Z interference, and an infrared singularity in the phase space; the lower electron mass provides a log-enhanced cross section.

[*]gmagill@perimeterinstitute.ca
[†]plestird@mcmaster.ca

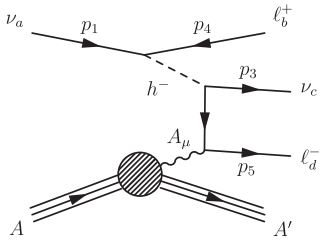

FIG. 1. Neutrino trident production of a charged Weyl lepton pair via a new charged scalar. There are three additional diagrams that can be obtained. The two charged leptons can be of different flavors. The connecting photon can interact with the nucleus (as shown above), or with individual nucleons.

Multiflavor configurations were not observable in CCFR or CHARM-II due to difficulties in tagging electron final states. The potential to view these NTP processes at future experiments allows for a rich landscape of signals [15,20,21]. In particular, it lends itself to the study of off-diagonal lepton flavor couplings, and these appear naturally for new charged scalars. In this work, we study how these new mixed-flavor observables compare with existing probes of charged scalar theories that preserve the SM's $\mathrm{SU}(2)\times\mathrm{U}(1)$; we assume no additional fermion or vector content. The case of neutral scalars probed via the diagonal $\nu\mu^+\mu^-$ final state has been considered for a phenomenologically motivated Lagrangian in Ref. [22].

We find that charged scalars are best probed by NTP in the case of universal flavor diagonal couplings. For these configurations, we find that mixed-flavor trident final states can give a 20%–50% increase in sensitivity to BSM couplings as compared to the traditional $\nu\mu^+\mu^-$-trident channel, and consequently outperforming bounds from the anomalous magnetic moment of the muon. When considering explicit UV completions (such as a Higgs triplet), we characterize the experimental improvements one should make in order for bounds from NTP to be competitive. Additional neutral and doubly charged scalar particles often appear in the context of UV models reproducing neutrino oscillation data, and these can introduce new, and more stringent, constraints.

The rest of the article is organized as follows: In Sec. II, we consider a general leptophilic charged scalar, how it contributes to trident, and its associated experimental backgrounds. For some benchmark choices in parameter space, we show the reach in sensitivity. In Sec. III, we explain how our general model can arise by giving an exhaustive classification of all leptophilic, renormalizable and $\mathrm{SU}(2)\times\mathrm{U}(1)$ invariant scalar extensions. We discuss specific realizations of these classifications in the literature and in Sec. IV, the phenomenological constraints surrounding them. We conclude with general remarks and potential applications in Sec. V.

II. CHARGED SCALAR MEDIATED TRIDENT PRODUCTION

A. Signal

We consider a singly electrically charged scalar h coupling to the lepton doublets,

$$\mathcal{L} \supset |\partial_\mu h|^2 + |\partial_\mu k|^2 - m_h^2|h|^2 + \sqrt{2}h_{ab}\nu^a\ell^b h + k_{ab}\ell^a\ell^b k + \text{c.c.} \tag{1}$$

The doubly electrically charged scalar k, which does not contribute to NTP, has been included to make connection with UV completions. The singly charged scalar contributes to NTP via diagrams like the one shown in Fig. 1 and results in the amplitude shown in Eq. (2). In the following, we use x^α and $y^\dagger_{\dot\alpha}$ to denote left- and right-handed initial states respectively, and y^α and $x^\dagger_{\dot\alpha}$ to denote right- and left-handed final states, following [23]. We assign the labels $\{1,2,3,4,5\} = \{\nu,\gamma,\nu',\ell^+,\ell^-\}$, and we use the mostly minus metric $\eta_{\mu\nu} = \{1,-1,-1,-1\}$.

In the context of the equivalent photon approximation [19,24], the matrix element for $\gamma\nu_a \to \ell_b^+\nu_c\ell_d^-$ can be summarized succinctly as

$$\begin{aligned}
\mathcal{M}_h &= -\frac{2h_{ab}h^*_{cd}e}{(P_1-P_4)^2-m_h^2}\left[\frac{\mathcal{A}_{14}}{q_d^2-m_d^2}+\frac{\mathcal{A}_{35}}{q_b^2-m_b^2}\right]\\
\mathcal{A}_{14} &= \{(x_1y_4)(x_3^\dagger\bar{\not{q}}_d\not{e}_2x_5^\dagger)+(x_1y_4)(x_3^\dagger\bar{\not{e}}_2y_5)m_d\}\\
\mathcal{A}_{35} &= \{(x_1\not{q}_b\bar{\not{e}}_2y_4)(x_3^\dagger x_5^\dagger)+(x_1\not{e}_2x_4^\dagger)(x_3^\dagger x_5^\dagger)m_b\}\\
q_b &\equiv P_2-P_4 \quad q_d \equiv P_2-P_5 \quad \not{\ell} = \ell\cdot\sigma \quad \bar{\not{\ell}} = \ell\cdot\bar\sigma,
\end{aligned} \tag{2}$$

where $\{a,b,c,d\}$ label lepton generations.

The above matrix element contains contributions from four different diagrams. Two contain mass insertions appearing in the second terms of $\mathcal{A}_{14}$ and $\mathcal{A}_{35}$. The two amplitudes correspond to the photon interacting with either the negatively or positively charged lepton. The following identities,

$$u=(x,y^\dagger)^T \quad v=(y,x^\dagger)^T \quad \bar u=(y,x^\dagger) \quad \bar v=(x,y^\dagger), \tag{3}$$

can be used to rewrite the amplitudes in terms of the Dirac spinors,

$$\begin{aligned}
\mathcal{A}_{14} &= \bar v_1 P_L v_4 \bar u_3 P_R(\not{q}_d+m_d)\not{\epsilon}_2 v_5,\\
\mathcal{A}_{35} &= \bar v_1 P_L(\not{q}_b+m_b)\not{\epsilon}_2 v_4 \bar u_3 P_R v_5.
\end{aligned} \tag{4}$$

As a check of our calculations, we used the symbolic manipulation language FORM [25] and compared our results to [20]. LEP searches rule out charged Higgs for $m_{h^\pm} \lesssim 100$ GeV on general grounds based solely on its electromagnetic interactions with the photon and Z boson

[26] and so we have ignored the four-momentum in the scalar's propagator. The full cross section is obtained from Eq. (2) by

$$\sigma_{N\nu} = \frac{Z^2\alpha}{\pi}\int_{m_{jk}^2}^{S}\frac{\mathrm{d}s}{s}\sigma_{\gamma\nu}(s)\int_{(s/2E_\nu)^2}^{\infty}\frac{\mathrm{d}Q^2}{Q^2}F^2(Q^2), \quad (5)$$

where $F(Q^2)$ above is the Woods-Saxon form factor [15,27].

For generic NTP final states the SM and BSM contributions can both be treated as real. The sign of the interference is dictated by the symmetry or antisymmetry of the couplings in Eq. (2), as well as the relative sign of the SM contribution. For a given NTP process, the presence of Z and/or W vector mediators induces an axial (C_A) and vector (C_V) coupling, upon which the matrix element depends linearly [15]. If the SM mediators are both W and Z bosons ($C_{V,A} > 0$), we find a positive relative sign. When the mediator is only a Z boson ($C_{V,A} < 0$), we get a negative sign. When the mediator is only a W boson ($C_{V,A} = 1$), we find a positive sign for $m_+ > m_-$ and a negative one when $m_+ < m_-$; this effect is related to subtle helicity properties [15]. For antisymmetric couplings $h_{e\mu} = -h_{\mu e}$ the new physics part of the matrix element carries an additional negative sign, while for the symmetric case ($h_{e\mu} = h_{\mu e}$), there is a positive sign. The final results for the sign of the interference terms are shown in Appendix C. For symmetric (antisymmetric) couplings, we have mostly constructive (destructive) interference.

B. Search strategy and backgrounds

Many flavor combinations for the incoming neutrino, outgoing neutrino, and charged leptons are possible. In deciding which reaction channel is ideally suited to one's purposes, two strategies should be considered. First, a channel with a relatively high SM contribution could be chosen, allowing for interference effects, which will be dominant in the limit of small coupling.[1] Neutrino beams are predominantly composed of ν_μ and so, in considering interference-driven signals, we typically consider incident ν_μ. Phase space considerations cause NTP rates to favor lighter lepton masses [15] and so we focus our analysis on final states with at least one electron, or positron. When considering the older experiments CCFR and CHARM-II we consider their reported observations of $\mu^+\mu^-$ production.

A complementary approach is to consider a production channel that is closed in the SM, but open in the case of new physics. To ensure low backgrounds, one needs to be able to control the flux of incident (anti)neutrinos. To see this consider $\nu_\mu \to e^-\mu^+\nu$ which is SM forbidden, but possible in the presence of BSM scalars. If, however, the beam was contaminated with $\bar{\nu}_\mu$, then the SM allowed $\bar{\nu}_\mu \to \mu^+ e^- \bar{\nu}_e$ would present a substantial background. DUNE has the capability to eliminate contamination with its neutrino horn. In contrast, SHiP has a much more complicated incident flux profile and cannot separate the neutrino and antineutrino fluxes.

The sensitivities we present in this paper are based on future experiments measuring rates consistent with the irreducible SM coherent NTP backgrounds. The details of the procedure are outlined in Appendix C, and numerical sensitivity equations are shown in Tables III and IV. A full simulation would have to be performed by the collaborations prior to their analysis, but we believe our analysis provides a good approximation. For simplicity, we focus the discussion on SHiP; however DUNE is also well equipped to tackle the same backgrounds.

The SHiP tau neutrino detector, modeled after the OPERA experiment [29], is based on emulsion cloud chambers (ECC) technology. The ECC is composed of a series of thin films interleaved with lead plates, followed by a muon spectrometer. A qualitatively similar setup to this was used in the CCFR experiment [14,30], which featured iron plates interleaved with liquid scintillators and drift chambers. The use of fine emulsion film layers will provide SHiP with more accurate track ID capabilities as compared to CCFR. That said, CCFR was able to observe a $\mu^+\mu^-$-trident rate of 37 events given a theoretical SM prediction of 45 events. They isolated their signal by collecting μ^- and μ^+ events and imposing cuts on the energy, angles, total invariant mass, hadronic activity, and vertex resolution. SHiP can implement similar cuts; however one caveat is that CCFR was dealing with much larger incoming neutrino average energies (~160 GeV) as compared with SHiP and DUNE providing it with an enhanced signal [15,19]. Since the bulk of these trident events are expected to come from SM processes, the kinematics of the outgoing pair of charged leptons is well captured in [20,21].

Consider a mixed-flavor $\ell_b^+\ell_d^-$ lepton pair search with a hadronic veto. Final e^+e^- states can arise from resonant π^0 production followed by a Dalitz decay where one of the photons is lost. For $\mu^-\mu^+$ final states, the dominant backgrounds are from $\nu_\mu A \to \mu^- Y X$, where Y represents either a pion, kaon, charm- or D-meson which decays to a final state involving μ^+ [31], as seen by NuTeV. Production of vector meson final states is also likely, but these can be distinguished from NTP since they deposit more hadronic energy and lead to a larger invariant mass for the lepton pair. The decay length of pions is on the order of a few meters, and therefore these backgrounds could also contribute to $\mu^- e^+$ mixed-flavor final states if the meson fakes a charged lepton before decaying. The fake rate suppression at SHiP is very competitive. In particular, for electron ID efficiencies greater than 80%, the pion contamination rate is roughly $\eta_{\pi\to e} \leq 0.5\%$ [32]. In the SHiP detector at

[1]This interference is not sensitive to the phases of the couplings, which can be expected on general grounds related to the arbitrary definitions of phases in h_{ab} [28].

the end of the decay chamber, pion contaminations of $\eta_{\pi\to\mu} = 0.1\%$ can be achieved for muon identification efficiencies of roughly 1. For e^- and μ^+ final states, it is difficult to imagine how this would be produced outside of NTP. One possibility is coherent pion production from a $\bar{\nu}_e$ incoming state and a negatively charged pion. This background is expected to be small for a number of reasons, owing to the differences in the ℓ^+ and ℓ^- energy spectrum, the much smaller lifetime flux of $\bar{\nu}_e$ at SHiP. Combinatorial backgrounds where one observes an electron and an antimuon from two unrelated processes could be eliminated by the micron vertex resolution available at both SHiP and DUNE.

C. Model independent results

In this section, we illustrate the sensitivity of mixed-flavor NTP to charged scalars. We also highlight how certain flavor configurations precluded in the SM give superior sensitivity to existing constraints. As an illustrative example, we consider the model described by Eq. (1) and assume that $h_{aa} = |h|$, and $h_{ab} = 0$ for $a \neq b$. As is eventually discussed in Appendix B, most of the strong existing constraints commonly considered for these types of models [10,11] drop out and NTP provides the dominant constraint, outperforming the $(g-2)$ for the muon. Our results are shown in Fig. 2.

We have forecasted the SM backgrounds at SHiP and DUNE using the rates presented in [15]. The best performing mode is the μ^+e^- channel at DUNE. *A priori*, the irreducible backgrounds to this process are $\bar{\nu}_\mu \to \mu^+\bar{\nu}_e e^-$ and $\nu_e \to \mu^+\nu_\mu e^-$. However, DUNE will have the ability to run in neutrino and antineutrino mode independently. This, coupled to the fact that the ν_e luminosities are low at this experiment, makes this channel a 0 irreducible background search. Hence, we can use this channel to investigate the interplay between 0 background and the lack of interference term in the cross section. We make the interesting

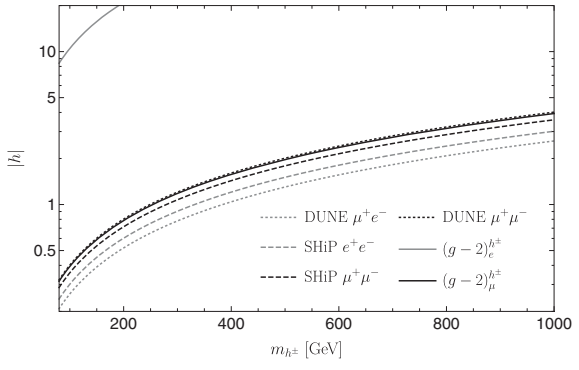

FIG. 2. Projected 90% C.L. sensitivities at DUNE and SHiP for a given pair of final state oppositely charged leptons, and competing constraints when allowing only $h_{ee} = h_{\mu\mu} = h_{\tau\tau} \neq 0$.

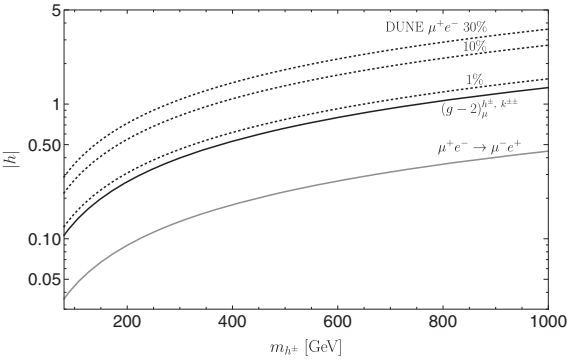

FIG. 3. Projected sensitivities of NTP assuming the SM prediction at DUNE has been measured to various precisions measured as a percentage of the SM cross section. We compare this to other constraints which now include a doubly charged scalar.

observation that the mixed-flavor final states in both experiments provide stronger constraints than the $\mu^+\mu^-$ states, while probing a Yukawa diagonal theory. In going from the muon final states to more general final states, the sensitivities to $|h|$ at DUNE are improved by 50% whereas at SHiP, they can be improved up to 20%.

We now show how NTP compares to other constraints when taking into account doubly charged scalars, assuming that $h_{ab} = k_{ab}$ in Eq. (2). This is analogous to the Higgs triplet (HT) model to be discussed later, without imposing the requirements of reproducing neutrino masses. The introduction of a doubly charged scalar $k^{\pm\pm}$ implies additional constraints from $\mu^+e^- \to \mu^-e^+$. A natural question to ask is the following: What improvements in sensitivity are required to make trident competitive with these stronger constraints? We assume that one could measure the NTP cross section to within a given percentage of the SM cross section, for various benchmark precisions. These results are presented in Fig. 3. As a reference, the 10% curve for DUNE's μ^+e^- channel corresponds roughly to the 90% C.L. bounds shown in Fig. 2. As can be seen, very high precision in the measured NTP cross section would be required to compete with the leading constraints on scalar couplings assuming $k^{\pm\pm}$.

III. EXTENSIONS ABOVE THE ELECTROWEAK SCALE

We now illustrate how the phenomenological charged scalar model from Sec. II can minimally arise while obeying all of the symmetries of the SM, with no additional fermion or vector matter content. The lepton sector's $\mathrm{SU}(2) \times \mathrm{U}(1)$ structure restricts possible scalar couplings that are relevant for NTP. The relevant leptonic fields are the SU(2) doublets and singlets denoted by

TABLE I. Classification of renormalizable lepton-scalar operators consistent with gauge invariance. The final column denotes the flavor symmetry ($\{ab\}$) or antisymmetry ($[ab]$) due to the SU(2) structure.

Field	U(1)	SU(2)	$\mathcal{L}_{\rm int}$	Couplings
$\mathcal{S}$	-2	1	$\mathcal{S}\ell_a^c\ell_b^c$	$s_{\{ab\}}$
$\mathcal{F}$	1	1	$\mathcal{F}\epsilon_{ij}L_a^{[i}L_b^{j]}$	$f_{[ab]}$
$\mathcal{D}$	$-\frac{1}{2}$	2	$\mathcal{D}_i L_a^i \ell_b^c$	d_{ab}
$\mathcal{T}$	1	3	$\mathcal{T}_{\{ij\}}L_a^{\{i}L_b^{j\}}$	$t_{\{ab\}}$

$$L_a^i = \begin{pmatrix} \nu_a \\ \ell_a \end{pmatrix}, \qquad \ell_a^c \tag{6}$$

respectively, where i labels the SU(2) index, and $a \in \{e, \mu, \tau\}$ labels the generations. All fields above are two-component left-handed spinors, with the spinor indices suppressed [i.e. $\ell_a^c = (\ell_a^c)_\alpha$]. To couple these fermions to a scalar via a renormalizable interaction we can consider at most two lepton fields and one scalar. The possibilities are given in Table I, where the lowercase letters represent generational coupling matrices, and the capital script letters are the scalar fields. In the order shown in the table, these are the symmetric singlet, antisymmetric singlet, doublet, and triplet models. The s_{ab} and t_{ab} couplings are symmetric in their indices. As for f_{ab}, the antisymmetry under the $i \leftrightarrow j$ forces f_{ab} to be antisymmetric under $a \leftrightarrow b$. The couplings for d_{ab} are unconstrained. A doubly charged scalar such as S cannot contribute to NTP at tree level, and we therefore focus on the fields $\mathcal{F}$, $\mathcal{D}$, and $\mathcal{T}$ for the purposes of NTP. The primary effects of the SU(2) symmetry are to

(i) Enforce a relation between couplings of the neutral, singly, and doubly charged scalars. This occurs for the triplet case and introduces additional constraints with which NTP must compete.

(ii) Generate flavor symmetries in the couplings which can lead to constructive or destructive interference.

To discuss specific implementations of the $\mathcal{D}$, $\mathcal{T}$ and $\mathcal{F}$ classifications, we respectively consider the Two-Higgs-doublet model, the HT model (also known as type-II seesaw), and the Zee-Babu (ZB) model. The full details of these models [10,11,28,33] are discussed in Appendix A. Here, we summarize the important features of the latter two theories. HT and ZB models both generate neutrino masses and feature a doubly charged scalar. In the ZB model, the couplings of leptons to the singly charged and doubly charged scalars are allowed to vary independently, whereas in the HT model, they are identical. In order to preserve the $\mathrm{SU}(2) \times \mathrm{U}(1)$ structure of the SM, the HT model contains in addition a neutral scalar which only couples to neutrinos. Without any extra model building, the neutral scalars considered in this paper cannot contribute to NTP in contrast to the models considered in Ref. [22]. To help make the connection with Sec. II, we show important coupling relations in Table II.

TABLE II. Relationships between type-II seesaw, Zee-Babu, and generic couplings h_{ab} and k_{ab}. Curly (square) braces mark the couplings' (anti)symmetry.

Scalar extension	$\sqrt{2}h_{ab}$	$h^\pm$	k_{ab}	$k^{\pm\pm}$
Zee-Babu	$2f_{[ab]}$	$\mathcal{F}$	$s_{\{ab\}}$	$\mathcal{S}$
Type-II seesaw	$\sqrt{2}t_{\{ab\}}$	$\Delta^\pm$	$t_{\{ab\}}$	$\Delta^{\pm\pm}$

IV. EXPLICIT UV COMPLETIONS

A. Singlet scalars

We consider the ZB model to demonstrate the effects of negative interference and the requirements of reproducing neutrino textures. Using Eqs. (A2) to (A4), we express all of the f_{ab} as a function of only $f_{e\mu}$ and Pontecorvo-Maki-Nakagawa-Sakata matrix data [34]. Note that due to the vanishing f_{aa} couplings of the ZB model, we are now probing nondiagonal couplings. We do this for both the normal and inverted hierarchies, and derive constraints on

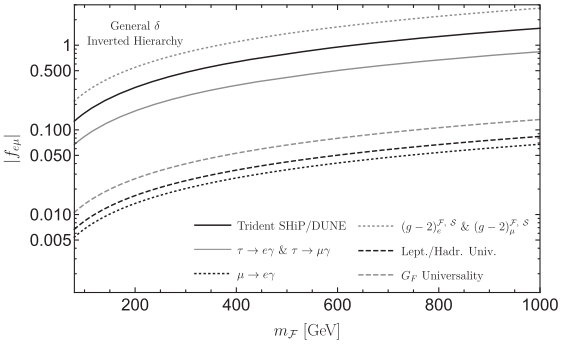

(a) Inverted Hierarchy

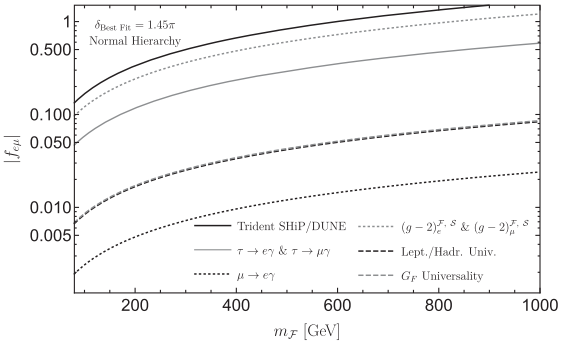

(b) Normal Hierarchy

FIG. 4. Sensitivities for $|f_{e\mu}|$ assuming the Zee-Babu model generates neutrino masses.

$|f_{e\mu}|$ as a function of $m_{\mathcal{F}}$ using the best performing mixed-flavor trident channels. For the normal hierarchy, we set the CP-violating phase δ to its best fit value. For the inverted hierarchy, the dependence on δ factors out, and so the ZB model's contribution to NTP is independent of δ. Our results are presented in Figs. 4(a) and 4(b).

B. Triplet scalars: Bounds from CCFR and CHARM-II

The $\nu_\mu \to \nu_i \mu^+ \mu^-$ final state was observed at the CCFR and CHARM-II experiments, and we can calculate experimental bounds on the triplet model using their data. Singlet scalars cannot be probed using this data due to the antisymmetry of the couplings f_{ab}.

CHARM-II had a neutrino beam of $\langle E_\nu \rangle \approx 20$ GeV [12,13] with a glass target ($Z = 11$) and the CCFR Collaboration had a neutrino beam of $\langle E_\nu \rangle \approx 160$ GeV using an iron target ($Z = 26$) [12,14]. The two experiments measured production cross sections of [12]

$$\begin{aligned} \sigma_{\text{CHARM-II}}/\sigma_{\text{SM}} &= 1.58 \pm 0.57, \\ \sigma_{\text{CCFR}}/\sigma_{\text{SM}} &= 0.82 \pm 0.28. \end{aligned} \tag{7}$$

Using CCFR as an example, we set bounds by demanding

$$\sigma_{\text{SM+Triplet}} \leq \sigma_{\text{SM}}(0.82 + 1.64 \times 0.28), \tag{8}$$

where 1.64 standard deviations encompasses 90% of a Gaussian likelihood function. For $m_{\mathcal{T}}$ in units of TeV,

$$|t_{\mu\mu}|^2 \left[\frac{26.38}{m_{\mathcal{T}}^2} + 1.59 \frac{|t_{e\mu}|^2 + |t_{\mu\mu}|^2 + |t_{\tau\mu}|^2}{m_{\mathcal{T}}^4} \right] \leq 691.36 \tag{9}$$

for CHARM-II, and

$$|t_{\mu\mu}|^2 \left[\frac{34.87}{m_{\mathcal{T}}^2} + 1.97 \frac{|t_{e\mu}|^2 + |t_{\mu\mu}|^2 + |t_{\tau\mu}|^2}{m_{\mathcal{T}}^4} \right] \leq 168.07 \tag{10}$$

for CCFR. Assuming $|t_{ab}| = |t|$, at 90% C.L. the two collaborations impose the following constraints:

$$\begin{aligned} |t| &\leq 3.10 \left(\frac{m_{\mathcal{T}}}{\text{TeV}} \right) \quad \text{CHARM-II}, \\ |t| &\leq 1.77 \left(\frac{m_{\mathcal{T}}}{\text{TeV}} \right) \quad \text{CCFR}. \end{aligned} \tag{11}$$

The stronger bounds from CCFR are a result of the fact that this experiment saw a deficit of events in comparison to the SM prediction and so the upper bound at 90% C.L. is lower than CHARM-II.

V. CONCLUSIONS AND OUTLOOK

We have investigated NTP as a tool for studying scalar extensions of the SM. We have considered SU(2) singlet, doublet, and triplet charged scalar extensions that couple to leptons, and concluded that triplet and singlet scalar can contribute appreciably to NTP.

In the case of triplet extensions we have found that NTP can serve as a complementary probe of the scalar sector at future experiments. In particular, for specific choices of model parameters in which LFV bounds vanish, NTP provides greater sensitivity than measurements of the anomalous magnetic moment. We found NTP to provide comparable sensitivity for charged singlet scalars and previous Z' models [12] in phenomenologically allowed mass ranges, despite their very different interaction nature. These prospects could be improved as the intensity frontier is pushed forward, and NTP may prove to be a valuable tool in the future. For generic choices of parameters, it is unlikely that NTP can compete with strong LFV constraints.

We have considered both the upcoming experiments SHiP and DUNE. The advantage of DUNE is its ability to isolate beams of ν_μ and $\bar{\nu}_\mu$ with high purity by using a magnetic horn. We have shown that this enables us to remove the irreducible background for certain processes, namely $\nu_\mu \to \mu^+ e^- \nu_i$, which has no SM contribution and is a viable production process in triplet models. This has the advantage of providing a clean signal, but results in a sensitivity that scales as $|h|^4$, in contrast to interference effects which can dominate for small coupling and scale as $|h|^2$. The lack of interference with the SM in these particular modes hinders sensitivity. For other channels, the relative phase between the SM and new physics contribution was found to be highly dependent on initial states, which had a tendency to cause destructive (constructive) interference in singlet (triplet) mediated NTP cross sections as can be seen in Appendix C.

The advantages provided by DUNE's nearly monoflavor beam must be balanced against its relatively low-Z detector (argon $Z = 18$) as compared to SHiP (lead $Z = 82$). Additionally DUNE uses a lower energy beam ($\langle E_\nu \rangle = 5$ GeV vs $\langle E_\nu \rangle = 20$ GeV) but compensates for this via a higher number of protons on target. In contrast to DUNE, SHiP's future lead-based detector provides an ideal setting to take advantage of the Z^2 coherent enhancement; however the lack of a neutrino horn, and the mutliflavor nature of the neutrino beam, suggests that searches at SHiP will have higher SM irreducible backgrounds.

Lastly, we have investigated representative UV models leading to the generic scalar extensions discussed above. In the ZB and HT models, extra particles and relations between couplings arise if the scalar sector is expected to produce empirically viable neutrino textures. The added constraints due to tree-level lepton flavor violating decays mediated by the doubly charged scalar and from the LHC are especially strong, and in some sense NTP is less important.

The influence of final states on the phase of the SM contribution may be of interest in future applications of NTP to new physics. This dependence is dictated not only

by the flavor combinations in the initial and final states, but also the relative sizes of the charged lepton masses. This final feature is a consequence of the chiral structure of the weak interaction [15]. The influence of these relative phases would be easy to miss and will play a crucial role in any future work that relies on interference with the SM. Although we have considered charged scalars which are already very constrained, we expect many of the qualitative features present in our analysis to be applicable to broader classes of model, in particular, the unique ability of mixed-flavor final states to control the presence or absence of constructive interference. Finally, we were able to identify final mixed-flavor states with no SM counterparts, thus removing irreducible backgrounds. Our results expand the reach of future neutrino experiments—such as DUNE, SHiP, and SBN—to physics beyond their main research program, both within and beyond the SM.

ACKNOWLEDGMENTS

We are very grateful to Itay Yavin and Maxim Pospelov for their continued guidance and for suggesting mixed-flavor trident production and its potential applicability to scalar models in future intensity frontier experiments. We thank Brian Shuve and Wolfgang Altmannshofer for feedback on the manuscript, as well as Cliff Burgess, Richard Hill, Stefania Gori, Chien-Yi Chen and Sarah Dawson for useful discussions. This research was supported in part by Perimeter Institute for Theoretical Physics. Research at Perimeter Institute is supported by the Government of Canada through the Department of Innovation, Science and Economic Development and by the Province of Ontario through the Ministry of Research and Innovation. This research was also supported by funds from the National Science and Engineering Research Council of Canada (NSERC), and the Ontario Graduate Scholarship (OGS) program.

APPENDIX A: EXPLICIT UV COMPLETIONS

In what follows we discuss popular implementations of each class of scalar models outlined above. The ZB model, used to radiatively generate neutrino masses, is a representative candidate for singlet scalars $\mathcal{F}$ (see Table I). Two-Higgs-doublet models (2HDM) have been considered extensively in the literature as an implementation of doublets $\mathcal{D}$ and we discuss neutrino trident production's ability to probe their couplings below. Finally the type-II seesaw mechanism (also known as HT models) for the generation of neutrino masses is discussed as the canonical example of a triplet model $\mathcal{T}$.

1. Two-Higgs-doublet models

2HDMs have been extensively studied [33,35,36]. In most implementations of a 2HDM there is mixing between the new BSM and SM Higgs doublets. This suggests that for the model to be technically natural couplings between the BSM charged doublet and leptons should be mass weighted to incorporate the influence of the mass-weighted SM Higgs field. In the SM, rates of NTP are log enhanced by infrared phase space effects which are controlled by the small masses of the charged leptons. If one were to consider NTP mediated by the charged component of a doublet scalar extension, this small mass infrared enhancement would compete directly against the mass-weighted Yukawa coupling suppression. Our explicit sensitivity calculations confirm that these competing effects make trident uncompetitive with existing constraints. We note that in the absence of mass-weighted couplings, NTP may be able to address this interesting region of parameter space; however this situation is technically unnatural due to radiative corrections from the Higgs boson—which induces corrections proportional to the SM Yukawa couplings—and would require a new physics mechanism to avoid fine-tuning.

2. Zee-Babu model

A popular implementation of the scalar singlet model is the ZB model [6,37,38]. The model features a singly charged scalar $\mathcal{F}$ that couples to the leptonic doublets, and a doubly charged scalar S which couples to the right-handed lepton singlets. The Yukawa sector of the Lagrangian can be written as

$$\begin{aligned}\mathcal{L}_{\mathrm{ZB}} &\supset f_{ab} L_a^i L_b^j \epsilon_{ij} \mathcal{F} + s_{ab} \ell_a^c \ell_b^c \mathcal{S} + \mathrm{H.c.} \\ &= 2 f_{ab} \nu_a \ell_b \mathcal{F} + s_{ab} \ell_a^c \ell_b^c \mathcal{S} + \mathrm{H.c.} \end{aligned} \tag{A1}$$

This model is typically considered in the context of radiatively generated neutrino masses. These first occur at two loops via diagrams such as the one shown in Fig. 5. Assuming the ZB model is fully responsible for the generation of neutrino masses, the mass matrix M can be expressed in terms of the ZB couplings f_{ab} and the SM Yukawa couplings Y via the relation $M \propto f Y s Y^T f^T$. The antisymmetric matrix f has odd dimensions and therefore its determinant will vanish by Jacobi's theorem. Since the neutrino mass matrix M contains f, its determinant will also vanish. This indicates that the smallest neutrino mass m_1 (m_3) will vanish in the case of the normal (inverted)

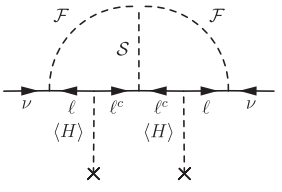

FIG. 5. Neutrino mass generation via the Zee-Babu model using two-component fermions with the direction of the arrows indicating chirality. [23].

hierarchy. The presence of a 0 mass mode [10,39] for the normal hierarchy implies

$$\frac{f_{e\tau}}{f_{\mu\tau}} = \tan\theta_{12}\frac{\cos\theta_{23}}{\cos\theta_{13}} + \tan\theta_{13}\sin\theta_{23}e^{-i\delta}, \quad \text{(A2)}$$

$$\frac{f_{e\mu}}{f_{\mu\tau}} = \tan\theta_{12}\frac{\sin\theta_{23}}{\cos\theta_{13}} - \tan\theta_{13}\cos\theta_{23}e^{-i\delta}, \quad \text{(A3)}$$

and for the inverted hierarchy yields

$$\frac{f_{e\tau}}{f_{\mu\tau}} = -\frac{\sin\theta_{23}}{\tan\theta_{13}}e^{-i\delta}, \qquad \frac{f_{e\mu}}{f_{\mu\tau}} = \frac{\cos\theta_{23}}{\tan\theta_{13}}e^{-i\delta}. \quad \text{(A4)}$$

These relations are used in the results of Sec. IV A, as they provide definite relations between the phases of the various couplings. A phase convention must be chosen, and a simple choice is $0 \leq f_{e\tau} \in \mathbb{R}$. Inspecting Eq. (A4) reveals that if the ZB model is responsible for the observed neutrino textures, and the hierarchy is determined to be inverted, then $0 > f_{e\mu} \in \mathbb{R}$, while $\mathrm{Arg} f_{\mu\tau} = \delta + \pi$. The case of the normal hierarchy is somewhat more involved; however two limits, namely $\delta = 0$ and $\delta = \pi$, result in all the couplings being real and positive by virtue of $\cos\theta_{23} \approx \sin\theta_{23}$ and $\tan\theta_{13} \ll \tan\theta_{12}$.

3. Type-II seesaw mechanism

One of the most popular triplet scalar extensions arises in the context of the seesaw mechanism for generating neutrino masses, specifically the so-called type-II seesaw or Higgs triplet model [40–42]. In this version, a triplet field with matrix representation

$$\mathcal{T} \equiv i\sigma_2 \cdot \Delta \equiv -\frac{1}{\sqrt{2}}\begin{pmatrix} -\Delta^0 & \Delta^+ \\ \Delta^+ & \sqrt{2}\Delta^{++} \end{pmatrix} \quad \text{(A5)}$$

is introduced into a symmetric lepton product via an interaction of the form $t_{ab}L_a^i\mathcal{T}_{ij}L_b^j$. After Δ^0 acquires a vacuum expectation value (VEV) v_T, we generate neutrino mass terms of the form $\mathrm{v}_T\nu_a\nu_b t_{ab}$. This model has been ruled out by measurements of the invisible width of the Z boson at LEP [42–44]. These bounds can be evaded by softly breaking the symmetry in the Lagrangian with the terms

$$-m_H^2H^2 + (\mu H^T i\sigma_2\Delta^\dagger H + \mathrm{H.c.}) + M_\Delta^2\mathrm{Tr}(\Delta^\dagger\Delta), \quad \text{(A6)}$$

where μ can be small to approximately retain the global symmetry. Minimizing this with respect to Δ^0 and setting $\langle\Delta^0\rangle = \mathrm{v}_T$ yields the equation

$$\mathrm{v}_T = \frac{\mathrm{v}_d^2\mu}{\sqrt{2}M_\Delta^2}, \quad \text{(A7)}$$

where v_d is the SM Higgs' VEV. Since the neutrino masses are given by $\mathrm{v}_T t_{ab}$, we can generate small masses in the limit where $M > \mathrm{v} \equiv \sqrt{\mathrm{v}_d^2 + \mathrm{v}_T^2} = 246$ GeV. As relevant for NTP, we have the Lagrangian

$$\mathcal{L}_{\mathrm{HT}} \supset -t_{ab}\left(\ell^a\Delta^+\sqrt{2}\nu^b + \ell^a\Delta^{++}\ell^b - \frac{\nu^a\Delta^0\nu^b}{\sqrt{2}}\right) + \mathrm{H.c.} \quad \text{(A8)}$$

The flavor symmetry of t_{ab} allows for flavor diagonal terms in the Lagrangian to be nonvanishing. This is in contrast to the singly charged couplings f_{ab} in the ZB model. The off-diagonal flavor couplings with $\Delta^\pm$ can be related to the ZB couplings as shown in Table II. The Δ^0 and $\Delta^{\pm\pm}$ scalars do not contribute to the NTP amplitudes, and so the trident exclusions we obtain on t_{ab} come only from $\Delta^\pm$ leptonic interactions. These must compete with other phenomenological considerations which can be mediated by the $\Delta^{\pm\pm}$ or Δ^0 fields. The propagating degrees of freedom of the scalar sector can, in principle, be different than the fields specified above [40]. However, ρ parameter constraints imply that the triplet VEV is at least 2 orders of magnitude smaller than the Higgs VEV [44]. Therefore the mixing is very small and we can think of $\Delta^\pm$ as being the physical mass eigenstate.

APPENDIX B: CONSTRAINTS

The addition of charged scalars to the SM leads to a variety of phenomenological consequences. In this section we discuss relevant constraints on the couplings involving the singly charged scalar (h_{ab}) and the couplings involving the doubly charged scalar (k_{ab}). The latter coupling does not play a role in NTP at tree level; however in the case of a triplet extension bounds on k_{ab} can be converted to constraints on h_{ab} since the two coupling matrices are related to one another as shown in Table II. Below we review existing probes of the scalar sector, which we compare with projected sensitivities using NTP as presented in Sec. IV A.

1. Anomalous magnetic moment measurements

Charged scalars can alter a particle's magnetic moment [45–47]. Additionally there is a long-standing discrepancy between the measured value of $(g-2)_\mu$ and the SM prediction [4]. As a result there is some ambiguity in the interpretation of this measurement as either a prediction of the BSM theory or as a constraint on its couplings. These bounds are the weakest of those presented in [10,11]. As shown in Sec. II for certain configurations of parameter space, NTP was capable of exceeding the sensitivity provided by this class of measurements. This is not surprising given NTP's competitive reach in the context of Z' models as outlined in [12].

2. Relative decay rates for μ and τ leptons

Another class of constraints can be obtained by using the relative size of various measured leptonic decay rates

$$\frac{\Gamma[\tau \to e/\mu + \text{inv.}]}{\Gamma[\mu \to e + \text{inv.}]} \quad \frac{\Gamma[\tau \to \mu + \text{inv.}]}{\Gamma[\tau \to e + \text{inv.}]} \tag{B1}$$

where inv. denotes invisible products (typically neutrinos). Measuring these quantities [48] effectively measures the deviation from unity of flavor ratios of weak couplings g_W^a/g_W^b for various flavors a and b. Models with charged scalars generically contribute to τ decays and so the measurements of Eq. (B1) can be translated as bounds on $||h_{i\tau}|^2 - |h_{ej}|^2|$ as a function of the mass of $h^\pm$ [10]. From the arguments of Ref. [10,28], a singly charged scalar would contribute to the decay $\mu \to e\nu\bar{\nu}$, but would not affect beta decay. Therefore by using data reported in Ref. [48] they were able to constrain $|h_{e\mu}|^2$, by considering a singly charged scalar's contribution to muon decay and noting that only final states with $e\nu_\mu\bar{\nu}_e$ would interfere with the SM amplitude.

The quoted constraint is $|h_{e\mu}|^2 < 0.014(\frac{m_h}{\text{TeV}})^2$ [10] after accounting for the normalizations shown in Table II. The full set of constraints as applied to the ZB model can be found in Tables II and III of Ref. [10], and model independent constraints for a singly charged scalar can be obtained by setting the doubly charged scalar's coupling to 0.

3. Loop-level LFV decays

LFV decays of the form $\ell_j \to \ell_i\gamma$ provide another tool to probe h_{ab}. This decay mode in the SM is extremely suppressed, and the observation of this LFV process would constitute strong evidence for new physics. Of particular interest is the decay mode $\mu \to e\gamma$ which provides the most stringent constraints on any of the couplings [49].

4. Tree-level LFV decays

In the case of triplet extensions where h_{ab} and k_{ab} are necessarily related (as shown in Table II) strong upper limits on certain decay modes [10], such as $\mu^- \to e^+e^-e^-$, already preclude the regions of parameter space trident is capable of probing. On some level these constraints may be evaded by choices related to the Majorana phases in the mass matrix [50,51]; however we have not included these subtleties in our analysis. For singlet scalar extensions k_{ab} and h_{ab} are independent and NTP does not need to compete with bounds related to tree-level LFV decays.

5. Implications of the LHC

When including doubly charged scalars, LHC constraints become very strong. There are analyses by both CMS and ATLAS [52,53] on doubly charged scalars decaying to same-sign dileptons which impose a model independent bound on the scalar mass of 200–400 GeV. A recast [54] of these same LHC searches extended the constraints on the mass by an additional 100 GeV by explicitly requiring a total nonzero lepton number in the final state (by considering final states of same-sign dileptons and gauge bosons). Furthermore, Ref. [55] has strengthened this bound by an additional 200 GeV by including recent data from the LHC's most recent 13 TeV run [55].

In Ref. [42], the authors showed that $h \to \gamma\gamma$ measurements at the LHC, the oblique T parameter limits and exclusions from LEP implies a lower bound on $m_{h^\pm}$ as a function of the triplet VEV v_T. The VEV enters in the generation of neutrino masses via the relation $M_{ab} = v_T t_{ab}$, as described in Appendix A 3. For example, $v_T \approx 1$ GeV implies $m_{h^\pm} \gtrsim 130$ GeV. This mass constraint gets stronger for lower values of v_T. Therefore, the Higgs triplet accounting for neutrino masses has very stringent limits. In the ZB model [10,11], the masses and couplings of the singly and doubly charged scalars can be independently tuned, subject to the constraint that the theory reproduce experimentally viable neutrino textures. There is therefore more flexibility in accommodating current data. In the scenario corresponding to an inverted neutrino mass hierarchy—among other assumptions—the constraints on the doubly charged scalar imply $m_{h^\pm} \geq 200$ GeV.

6. Neutrino masses

When considering neutrino masses, there are other sources of constraints that arise in addition to lepton flavor violation. The neutrino mass mixing matrix is related to the scalar triplet's couplings by $m_{ab} = v_T t_{ab}$. Hence, the sensitivity one must achieve in t_{ab} to probe the neutrino mass sector scales inversely with the VEV of the Higgs triplet. This favors using NTP to probe lower values of v_T. However, as was discussed in Appendix A 3, this implies a larger $m_{h^\pm}$. Coupled with recent cosmological bounds on the sum of neutrino masses [56], this makes NTP uncompetitive; we have confirmed this fact numerically.

APPENDIX C: PROJECTED SENSITIVITIES

Given the posterior distribution $P(\theta|\vec{x})$, we can define a 90% C.L. interval [57]. Making use of Bayes' theorem, we can express the posterior probability in terms of a Poisson likelihood, a prior—which is a step function in the signal event rate—and a normalization. The mean of the Poisson distribution is given by $\theta' = B + S$, where B is the background prediction and S is the signal events. Since there is no data $\vec{x}$, we assume that the future experiments will have observed the predicted number of background events. Collecting everything, we have

$$1-\alpha=\int_{-\infty}^{\theta_{\rm up}}P(\theta'|\vec{x})d\theta'=1-\frac{\Gamma(1+B,B+\theta_{\rm up})}{\Gamma(1+B,B)}, \tag{C1}$$

and solve for $\theta_{\rm up}$ given $\alpha=0.1$. Setting $B+\theta_{\rm up}=N_{\rm SM+NP}$, we can set 90% C.L. bounds on the couplings as a function of the masses of the new charged scalars. At SHiP, we take into account backgrounds from incoming ν and $\bar{\nu}$ whereas at DUNE, we consider only incoming ν. For both collaborations, the signal dependence takes only into account incoming ν. The mass of the new scalar is assumed to be in TeV.

1. SU(2) singlet scalar extensions

TABLE III. Projected 90% C.L. sensitivity for a variety of NTP processes mediated by an SU(2) singlet scalar with unit charge at both SHiP and DUNE.

Final state	SHiP	DUNE near detector
$e^+\mu^-$	$17.78\geq\frac{0.62\lvert f_{\mu e}\rvert^2(\lvert f_{e\mu}\rvert^2+\lvert f_{\tau\mu}\rvert^2)}{m_{\mathcal{F}}^4}-\frac{14.47\lvert f_{e\mu}\rvert^2}{m_{\mathcal{F}}^2}$	$15.53\geq\frac{0.53\lvert f_{\mu e}\rvert^2(\lvert f_{e\mu}\rvert^2+\lvert f_{\tau\mu}\rvert^2)}{m_{\mathcal{F}}^4}-\frac{12.66\lvert f_{e\mu}\rvert^2}{m_{\mathcal{F}}^2}$
e^+e^-	$16.82\geq-\frac{5.56\lvert f_{e\mu}\rvert^2}{m_{\mathcal{F}}^2}+\frac{1.66\lvert f_{\mu e}\rvert^2(\lvert f_{\mu e}\rvert^2+\lvert f_{\tau e}\rvert^2)}{m_{\mathcal{F}}^4}$	$9.38\geq-\frac{4.48\lvert f_{e\mu}\rvert^2}{m_{\mathcal{F}}^2}+\frac{1.35\lvert f_{\mu e}\rvert^2(\lvert f_{\mu e}\rvert^2+\lvert f_{\tau e}\rvert^2)}{m_{\mathcal{F}}^4}$

(C2)

2. SU(2) triplet scalar extensions

TABLE IV. Projected 90% C.L. sensitivity for a variety of NTP processes mediated by the singly charged component of an SU(2) triplet scalar field at both SHiP and DUNE.

Final state	SHiP
$e^+\mu^-$	$17.78\geq\frac{0.04\lvert t_{ee}\rvert^2(\lvert t_{e\mu}\rvert^2+\lvert t_{\mu\mu}\rvert^2+\lvert t_{\tau\mu}\rvert^2)}{m_{\mathcal{T}}^4}+\frac{0.16\lvert t_{\mu e}\rvert^2(\lvert t_{e\mu}\rvert^2+\lvert t_{\mu\mu}\rvert^2+\lvert t_{\tau\mu}\rvert^2)}{m_{\mathcal{T}}^4}+\frac{7.24\lvert t_{e\mu}\rvert^2}{m_{\mathcal{T}}^2}$
e^+e^-	$16.82\geq\frac{0.07\lvert t_{ee}\rvert^2(\lvert t_{ee}\rvert^2+\lvert t_{\mu e}\rvert^2+\lvert t_{\tau e}\rvert^2)}{m_{\mathcal{T}}^4}+\frac{0.42\lvert t_{\mu e}\rvert^2(\lvert t_{ee}\rvert^2+\lvert t_{\mu e}\rvert^2+\lvert t_{\tau e}\rvert^2)}{m_{\mathcal{T}}^4}+\frac{1.23\lvert t_{ee}\rvert^2}{m_{\mathcal{T}}^2}-\frac{2.78\lvert t_{e\mu}\rvert^2}{m_{\mathcal{T}}^2}$
$\mu^+\mu^-$	$6.43\geq\frac{0.01\lvert t_{e\mu}\rvert^2(\lvert t_{e\mu}\rvert^2+\lvert t_{\mu\mu}\rvert^2+\lvert t_{\tau\mu}\rvert^2)}{m_{\mathcal{T}}^4}+\frac{0.02\lvert t_{\mu\mu}\rvert^2(\lvert t_{e\mu}\rvert^2+\lvert t_{\mu\mu}\rvert^2+\lvert t_{\tau\mu}\rvert^2)}{m_{\mathcal{T}}^4}-\frac{0.04\lvert t_{e\mu}\rvert^2}{m_{\mathcal{T}}^2}+\frac{0.28\lvert t_{\mu\mu}\rvert^2}{m_{\mathcal{T}}^2}$
μ^+e^-	$11.65\geq\frac{0.02\lvert t_{e\mu}\rvert^2(\lvert t_{ee}\rvert^2+\lvert t_{\mu e}\rvert^2+\lvert t_{\tau e}\rvert^2)}{m_{\mathcal{T}}^4}+\frac{0.07\lvert t_{\mu\mu}\rvert^2(\lvert t_{ee}\rvert^2+\lvert t_{\mu e}\rvert^2+\lvert t_{\tau e}\rvert^2)}{m_{\mathcal{T}}^4}-\frac{0.38\lvert t_{e\mu}\rvert\lvert t_{\mu e}\rvert}{m_{\mathcal{T}}^2}$

Final state	DUNE near detector
$e^+\mu^-$	$15.53\geq\frac{0.13\lvert t_{\mu e}\rvert^2(\lvert t_{e\mu}\rvert^2+\lvert t_{\mu\mu}\rvert^2+\lvert t_{\tau\mu}\rvert^2)}{m_{\mathcal{T}}^4}+\frac{6.33\lvert t_{e\mu}\rvert^2}{m_{\mathcal{T}}^2}$
e^+e^-	$9.38\geq\frac{0.34\lvert t_{\mu e}\rvert^2(\lvert t_{ee}\rvert^2+\lvert t_{\mu e}\rvert^2+\lvert t_{\tau e}\rvert^2)}{m_{\mathcal{T}}^4}+\frac{0.04\lvert t_{ee}\rvert^2}{m_{\mathcal{T}}^2}-\frac{2.24\lvert t_{e\mu}\rvert^2}{m_{\mathcal{T}}^2}$
$\mu^+\mu^-$	$3.9\geq\frac{0.01\lvert t_{\mu\mu}\rvert^2(\lvert t_{e\mu}\rvert^2+\lvert t_{\mu\mu}\rvert^2+\lvert t_{\tau\mu}\rvert^2)}{m_{\mathcal{T}}^4}+\frac{0.12\lvert t_{\mu\mu}\rvert^2}{m_{\mathcal{T}}^2}$
μ^+e^-	$2.59\geq\frac{0.06\lvert t_{\mu\mu}\rvert^2(\lvert t_{ee}\rvert^2+\lvert t_{\mu e}\rvert^2+\lvert t_{\tau e}\rvert^2)}{m_{\mathcal{T}}^4}-\frac{0.01\lvert t_{e\mu}\rvert\lvert t_{\mu e}\rvert}{m_{\mathcal{T}}^2}$

(C3)

[1] R. Pohl, R. Gilman, G. A. Miller, and K. Pachucki, Muonic hydrogen and the proton radius puzzle, Annu. Rev. Nucl. Part. Sci. **63**, 175 (2013).

[2] R. J. Hill, in Proceedings, 12th Conference on Quark Confinement and the Hadron Spectrum (Confinement XII), Thessaloniki, Greece [*EPJ Web Conf.* 137, 01023 (2017).

[3] F. Jegerlehner and A. Nyffeler, The muon g-2, Phys. Rep. **477**, 1 (2009).

[4] G. W. Bennett *et al.* (Muon g-2 Collaboration), Final report of the muon E821 anomalous magnetic moment measurement at BNL, Phys. Rev. D **73**, 072003 (2006).

[5] A. Aguilar-Arevalo *et al.* (LSND Collaboration), Evidence for neutrino oscillations from the observation of antineutrino (electron) appearance in a antineutrino (muon) beam, Phys. Rev. D **64**, 112007 (2001).
[6] K. S. Babu, Model of calculable Majorana neutrino masses, Phys. Lett. B **203**, 132 (1988).
[7] K. S. Babu and S. Pakvasa, Lepton number violating muon decay and the LSND neutrino anomaly, arXiv:hep-ph/0204236.
[8] M. Lindner, M. Platscher, and F. S. Queiroz, A call for new physics: The muon anomalous magnetic moment and lepton flavor violation, arXiv:1610.06587 [Phys. Rep. (to be published)].
[9] Y.-S. Liu, D. McKeen, and G. A. Miller, Electrophobic Scalar Boson and Muonic Puzzles, Phys. Rev. Lett. **117**, 101801 (2016).
[10] J. Herrero-Garcia, M. Nebot, N. Rius, and A. Santamaria, The Zee-Babu model revisited in the light of new data, Nucl. Phys. **B885**, 542 (2014).
[11] P. S. B. Dev, C. M. Vila, and W. Rodejohann, Naturalness in testable type II seesaw scenarios, Nucl. Phys. **B921**, 436 (2017).
[12] W. Altmannshofer, S. Gori, M. Pospelov, and I. Yavin, Neutrino Trident Production: A Powerful Probe of New Physics with Neutrino Beams, Phys. Rev. Lett. **113**, 091801 (2014).
[13] D. Geiregat *et al.* (CHARM-II Collaboration), First observation of neutrino trident production, Phys. Lett. B **245**, 271 (1990).
[14] S. R. Mishra *et al.* (CCFR Collaboration), Neutrino Tridents and W Z Interference, Phys. Rev. Lett. **66**, 3117 (1991).
[15] G. Magill and R. Plestid, Neutrino trident production at the intensity frontier, Phys. Rev. D **95**, 073004 (2017).
[16] R. Acciarri *et al.* (DUNE Collaboration), Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE), arXiv:1512.06148.
[17] M. Anelli *et al.* (SHiP Collaboration), A facility to Search for Hidden Particles (SHiP) at the CERN SPS, arXiv:1504.04956.
[18] H. Chen *et al.* (MicroBooNE Collaboration), *Proposal for a New Experiment Using the Booster and NuMI Neutrino Beamlines: MicroBooNE* (2007).
[19] R. Belusevic and J. Smith, W-z interference in ν-nucleus scattering, Phys. Rev. D **37**, 2419 (1988).
[20] J. Lovseth and M. Radomiski, Kinematical distributions of neutrino-produced lepton triplets, Phys. Rev. D **3**, 2686 (1971).
[21] R. W. Brown, R. H. Hobbs, J. Smith, and N. Stanko, Intermediate boson. iii. Virtual-boson effects in neutrino trident production, Phys. Rev. D **6**, 3273 (1972).
[22] S.-F. Ge, M. Lindner, and W. Rodejohann, Atmospheric trident production for probing new physics, Phys. Lett. B **772**, 164 (2017).
[23] H. K. Dreiner, H. E. Haber, and S. P. Martin, Two-component spinor techniques and Feynman rules for quantum field theory and supersymmetry, Phys. Rep. **494**, 1 (2010).
[24] V. M. Budnev, I. F. Ginzburg, G. V. Meledin, and V. G. Serbo, The two photon particle production mechanism. Physical problems. Applications. Equivalent photon approximation, Phys. Rep. **15**, 181 (1975).
[25] J. A. M. Vermaseren, New features of FORM, arXiv:math-ph/0010025.
[26] G. Abbiendi *et al.* (LEP, DELPHI, OPAL, ALEPH, L3 Collaborations), Search for charged Higgs bosons: Combined results using LEP data, Eur. Phys. J. C **73**, 2463 (2013).
[27] U. D. Jentschura and V. G. Serbo, Nuclear form factor, validity of the equivalent photon approximation and Coulomb corrections to muon pair production in photon-nucleus and nucleus-nucleus collisions, Eur. Phys. J. C **64**, 309 (2009).
[28] M. Nebot, J. F. Oliver, D. Palao, and A. Santamaria, Prospects for the Zee-Babu Model at the CERN LHC and low energy experiments, Phys. Rev. D **77**, 093013 (2008).
[29] N. Agafonova *et al.* (OPERA Collaboration), Discovery of τ Neutrino Appearance in the CNGS Neutrino Beam with the OPERA Experiment, Phys. Rev. Lett. **115**, 121802 (2015).
[30] B. J. King *et al.*, Measuring muon momenta with the CCFR neutrino detector, Nucl. Instrum. Methods Phys. Res., Sect. A **302**, 254 (1991).
[31] T. Adams *et al.* (NuTeV Collaboration), Evidence for diffractive charm production in muon-neutrino Fe and antimuon-neutrino Fe scattering at the Tevatron, Phys. Rev. D **61**, 092001 (2000).
[32] A. Buonaura (SHiP Collaboration), *The SHiP experiment and its detector for neutrino physics* (2016).
[33] J. Cao, P. Wan, L. Wu, and J. M. Yang, Lepton-specific two-Higgs doublet model: Experimental constraints and implication on Higgs phenomenology, Phys. Rev. D **80**, 071701 (2009).
[34] I. Esteban, M. C. Gonzalez-Garcia, M. Maltoni, I. Martinez-Soler, and T. Schwetz, Updated fit to three neutrino mixing: Exploring the accelerator-reactor complementarity, J. High Energy Phys. 17 (2017) 87.
[35] T. Abe, R. Sato, and K. Yagyu, Lepton-specific two-Higgs doublet model as a solution of muon g-2 anomaly, J. High Energy Phys. 07 (2015) 064.
[36] G. C. Branco, P. M. Ferreira, L. Lavoura, M. N. Rebelo, M. Sher, and J. P. Silva, Theory and phenomenology of two-Higgs-doublet models, Phys. Rep. **516**, 1 (2012).
[37] T. P. Cheng and L.-F. Li, Neutrino masses, mixings and oscillations in $\mathrm{SU}(2) \times \mathrm{U}(1)$ models of electroweak interactions, Phys. Rev. D **22**, 2860 (1980).
[38] A. Zee, Charged scalar field and quantum number violations, Phys. Lett. **161B**, 141 (1985).
[39] J. Herrero-Garcia, M. Nebot, N. Rius, and A. Santamaria, in Proceedings, 37th International Conference on High Energy Physics (ICHEP 2014), Valencia, Spain, 2014 [*Nucl. Part. Phys. Proc.* 273–275, 1678 (2016)].
[40] P. F. Perez, T. Han, G.-y. Huang, T. Li, and K. Wang, Neutrino masses and the CERN LHC: Testing type-II seesaw, Phys. Rev. D **78**, 015018 (2008).
[41] H. Sugiyama, in *Proceedings, 1st Toyama International Workshop on Higgs as a Probe of New Physics 2013 (HPNP2013), Toyama, Japan, 2013* (2013).
[42] D. Das and A. Santamaria, Updated scalar sector constraints in the Higgs triplet model, Phys. Rev. D **94**, 015015 (2016).

[43] J. C. Montero, C. A. de S. Pires, and V. Pleitez, Spontaneous breaking of a global symmetry in a 331 model, Phys. Rev. D **60**, 115003 (1999).

[44] P. Langacker, *The Standard Model and Beyond, Series in High Energy Physics, Cosmology, and Gravitation* (Taylor and Francis, Boca Raton, FL, 2010).

[45] Y.-L. Wu and Y.-F. Zhou, Muon anomalous magnetic moment in the standard model with two Higgs doublets, Phys. Rev. D **64**, 115018 (2001).

[46] D. A. Dicus, H.-J. He, and J. N. Ng, Neutrino-Lepton Masses, Zee Scalars, and Muon g-2, Phys. Rev. Lett. **87**, 111803 (2001).

[47] C. A. de S. Pires and P. S. Rodrigues da Silva, Scalar scenarios contributing to (g-2)(muon) with enhanced Yukawa couplings, Phys. Rev. D **64**, 117701 (2001).

[48] A. Pich, Precision tau physics, Prog. Part. Nucl. Phys. **75**, 41 (2014).

[49] A. M. Baldini *et al.* (MEG Collaboration), Search for the lepton flavor violating decay $\mu^+ \to e^+\gamma$ with the full data set of the MEG experiment, Eur. Phys. J. C **76**, 434 (2016).

[50] A. Merle and W. Rodejohann, The Elements of the neutrino mass matrix: Allowed ranges and implications of texture zeros, Phys. Rev. D **73**, 073012 (2006).

[51] W. Grimus and P. O. Ludl, Correlations of the elements of the neutrino mass matrix, J. High Energy Phys. 12 (2012) 117.

[52] S. Chatrchyan *et al.* (CMS Collaboration), A search for a doubly charged Higgs boson in pp collisions at $\sqrt{s} = 7$ TeV, Eur. Phys. J. C **72**, 2189 (2012).

[53] G. Aad *et al.* (ATLAS Collaboration), Search for doubly charged Higgs bosons in like-sign dilepton final states at $\sqrt{s} = 7$ TeV with the ATLAS detector, Eur. Phys. J. C **72**, 2244 (2012).

[54] F. del Aguila and M. Chala, LHC bounds on lepton number violation mediated by doubly and singly charged scalars, J. High Energy Phys. 03 (**2014**) 027.

[55] J. Alcaide, M. Chala, and A. Santamaria, LHC signals of radiatively induced neutrino masses and implications for the Zee-Babu model, arXiv:1710.05885.

[56] F. Couchot, S. Henrot-Versillé, O. Perdereau, S. Plaszczynski, B. Rouillé d'Orfeuil, M. Spinelli, and M. Tristram, Cosmological constraints on the neutrino mass including systematic uncertainties, Astron. Astrophys. **606**, A104 (2017).

[57] K. A. Olive *et al.* (Particle Data Group), Review of particle physics, Chin. Phys. C **38**, 090001 (2014).

Chapter 7

Probing Compressed States at the LHC

7.1 Preface

This paper is currently in preparation. A draft has been written by Gabriel Magill for simulation results he obtained 2 years ago. The simulation and analysis that was done for the draft was based on a 8 TeV CMS search in 2014 [94]. As of recently, an updated version of the search was performed at 13 TeV by the ATLAS collaboration in 2017 [95]. The cuts imposed in the 2017 search are sufficiently different and more stringent than the ones from the 2014 search. Therefore, updating the simulation and the analysis in this draft to reflect the 2017 search would be necessary for a publication. As a result of not having gone through peer review processes, the material contained in this draft should be considered as *Preliminary Information*!

That said, the draft that was prepared for submission required substantial work and computing time, and is included here. The results obtained are in themselves interesting, as they identify a new topology that has thus far never been considered at the LHC. In the interest of maximizing the reach of the LHC to models that may be hidden from conventional LHC searches, the author of this thesis believes it is important to include this draft in its current form in the

thesis. The new topology we consider are soft leptons, missing energy and many jets. We motivate this new topology by developing a search for new colored doublet scalars in a region of parameter space that is typically very challenging to probe, namely when the mass splitting between the states of the doublet is small with respect to the weak scale.

Contribution of author (reproduced from Declaration of Authorship): Gabriel Magill performed all of the signal and background simulations, and developed most of the analysis code for the recast and projected sensitivity study (with help from Brian Shuve). Gabriel also contributed significantly to the writing of the draft. Jonathan calculated the decay length and branching ratio for the new colored scalars, and simulated the $2j$ QCD background process.

7.2 Paper

Prepared for submission to JHEP

Probing Compressed States at the LHC

G. Magill[a,b] **B. Shuve**[c] **J. Kozaczuk**[d] **E. Izaguirre**[e]

[a] *Perimeter Institute for Theoretical Physics, 31 Caroline St. N., Waterloo, Ontario N2L 2Y5, Canada*

[b] *Department of Physics & Astronomy, McMaster University, 1280 Main St. W., Hamilton, Ontario L8S 4M1, Canada*

[c] *Harvey Mudd College, 301 Platt Blvd., Claremont, CA 91711, USA*

[d] *Amherst Center for Fundamental Interactions, Department of Physics, University of Massachusetts, Amherst, MA 01003, USA*

[e] *Brookhaven National Laboratory, Upton, NY 11973, USA*

Abstract: In this paper, we present a comprehensive study detailing the gains in sensitivity from considering soft leptons and low missing energy signal regions in multijet collider searches. We demonstrate the power of these signal regions by considering colored scalar doublet models that participate in cascade decays. When no requirements are made on leptons, these signal events need to compete with a very large QCD background, and can be hidden from conventional searches. When the doublet mass splitting is small, we show that imposing simple kinematic upper bounds in the associated lepton channel can significantly boost sensitivity, without the need for a dedicated analysis. Additional handles on theories with a compressed mass splitting is of interest as these regions in parameter space are typically difficult to detect at hadron colliders. We demonstrate how our results are robust to changes in the hadronic part of the analysis, and how they depend on b-tag requirements.

Contents

1 Introduction

The main objective of the Large Hadron Collider (LHC) is to study physics at the electroweak scale. So far, their searches have shown very impressive agreement with the predictions of the Standard Model (SM), but not much insofar as new physics is concerned. In light of this, it is important to revisit the assumptions that are made in the searches as they might be hindering sensitivity to other physics scenarios. One feature that is very common across most searches is the requirement (or more generally the lack of veto) of large missing energy ($\not{E}_T$) signatures. The interest in large $\not{E}_T$ has a solid theoretical basis, since it often

arises from solutions of dark matter and the hierarchy problem in which stable new states escape detection.

It is however conceivable that there may exist new physics within our grasp that does not produce large $\not{E}_T$. This possibility has started to gain popularity. For instance, in the context of supersymmetry, the stability of the lightest new states is controlled by a discrete symmetry called R-parity. Although often imposed, it is not strictly necessary in order to solve the hierarchy problem. Various LHC searches have already considered topologies in which this symmetry is broken [1], leading to multijet final states. Similar ideas have also been put forward in the context of stealth supersymmetry [2, 3].

Building upon these ideas, we consider how sensitivity to novel signals can be achieved when explicitly imposing low $\not{E}_T$ cuts in multijet topologies, and exploring the presence of leptons as a QCD background reduction tool. This phenomenology can be achieved quite generically if the new physics produces soft off-shell W bosons. When the W decays leptonically, both the neutrino and the lepton's transverse momentum (p_T) will favor low values. A simple method of producing soft W bosons is via cascade decays of SU(2) doublets, in which the states of the scalar doublet have a small (compressed) mass splitting. An illustration of this process is shown in fig. 1. This is a very interesting scenario to study, since compressed electroweak states are often invisible to conventional searches and require dedicated analyses. We will show that including very simple additional lepton and $\not{E}_T$ signal regions on top of conventional searches can already be quite powerful in targeting these hidden states, without the need for a separate analysis.

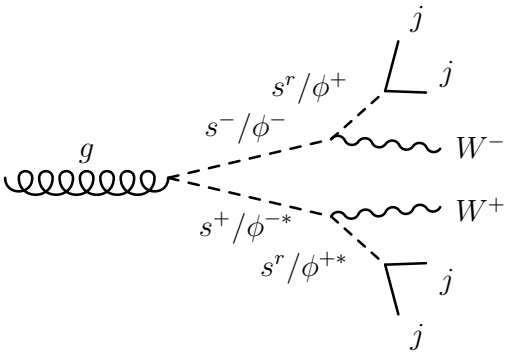

Figure 1: Illustrative Feynman diagram depicting the cascade decays of interest for the octet/triplet scalar doublet models. When the mass difference between the heavy and light partners of the doublet is less than m_W, the off-shell W boson will produce soft leptons and little missing energy.

As an explicit realization of electroweak states with a small mass splitting, we consider colored scalar doublet extensions of the Higgs sector. Having the scalars charged under SU(3)ensures that they will be copiously produced at hadron colliders. It has been shown

by Manohar and Wise [4] that there are only a handful of allowed renormalizable colored scalar extensions that are compatible with minimal flavor violation and hence avoid flavor changing neutral currents. Under SU(3)×SU(2)×U(1), one possibility is a $(\mathbf{8},\mathbf{2})_1$ octet scalar. If we allow ourselves to consider non-renormalizable theories, we can have a $(\mathbf{3},\bar{\mathbf{2}})_Y$ triplet scalar, where Y is determined by hypercharge conservation. Lastly, if the octet scalar couples to heavy vector-like fermions, integrating these out these heavy states can give rise to effective higher dimensional operators. We will refer to this last scenario as the octet-UV model.

These models have already been previously considered [5, 6]. In our analysis we use these models as a way of quantifying the gains in sensitivity when imposing additional soft lepton and low $\not{E}_T$ signal region cuts. We also examine the interplay between these cuts and the number of b-tags in the final state. Benchmarking these various channels is very useful for two reasons. Firstly, large systematic uncertainties associated in one channel can be averted by focusing instead on another channel with comparable sensitivity. Secondly, powerful b-tag requirements can be replaced with soft lepton requirements for signals that don't produce many b-jets.

The structure of the paper is as follows. In section 2, we highlight important features of our three colored scalar doublet models, and motivate searching for them via cascade decays. In section 4, we introduce our proposals for additional new signal regions on the lepton p_T and event $\not{E}_T$. This is done by simulating signal and background final state kinematics in a full projected sensitivity analysis for the 0, 1 and 2 lepton channels, and comparing to a recasted single lepton search. In section 5, we show sensitivity results for the projected analysis and the recasted search. We then summarize our findings in section 6 and discuss possible extensions of our ideas.

2 Representative Models

In order to test our proposal on the importance of soft leptons and low $\not{E}_T$ as a powerful background reduction tool in a multijet topology, we consider 3 different scalar SU(2) doublet models:

- Colored triplet scalars, discussed in section 2.1 and appendix A.1.
- Colored octet scalars, discussed in section 2.2 and appendix A.2.
- Colored octet-UV scalars, discussed in section 2.3.

From these models, our main signature will arise from an electroweak cascade decay of a heavy scalar partner to a light scalar partner and a W boson, as explained in section 3.2. In order for the W boson to produce a soft lepton and neutrino, we need the mass splitting between the heavy and light scalar states to be less than m_W. We illustrate how this can naturally happen using the octet model for concreteness. Following the derivation in

appendix A.2, the mass splitting between the two octet states ($S^{\pm}$ and S^r) is given as

$$|m_{S^\pm} - m_{S^r}| = \frac{v^2(\lambda_2 + 2\lambda_3)}{4(m_{S^\pm} + m_{S^r})}, \tag{2.1}$$

where v is the Higgs vacuum expectation value and λ_i are various $\mathcal{O}(1)$ coefficients in front of quartic scalar interaction terms. The smallest values of the light scalars we will consider in this paper range between $400 - 600$GeV. Therefore numerically, taking $v = 246$GeV and $m_{S^\pm} + m_{S^r} \geq 800$GeV, eq. (2.1) becomes

$$|m_{S^\pm} - m_{S^r}| \leq 19\text{GeV}(\lambda_2 + 2\lambda_3). \tag{2.2}$$

Choosing for example $\lambda_2 = \lambda_3 = 0.2$, we naturally obtain small mass splittings on the order of 10GeV. We now give a brief overview of each of these models and highlight important features.

2.1 Colored Triplets

We will now summarize the main features of the colored triplet model, leaving the full details in appendix A.1. The model consists of two scalar doublets Φ and Φ', both in the $(\mathbf{3}, \bar{\mathbf{2}})_Y$ representation, with $Y(\Phi) = +1/6$ and $Y(\Phi') = -5/6$. The doublet components of Φ are denoted as (φ^-, φ^+), although the specific electric charges depend on the value of Y. As we will make more explicit in section 3.1, we focus the discussion on the Yukawa interactions of φ^+ (which we take to be lighter than φ^-) with two down type quarks of flavor k and l. Since top quarks are not involved, in this discussion we will assume to a very good approximation that $m_{\varphi^+} \gg m_k, m_l$. The subset of the Lagrangian involving φ^+ is given in terms of Dirac and Weyl spinors [7] as

$$\begin{aligned} \mathcal{L} \supset & \frac{-i}{2\Lambda_{kl}}\varphi^+ \left[\bar{d^c}_k(m_k + m_l) - \gamma^5(m_k - m_l)d_l\right] \\ = & \frac{-i}{\Lambda_{kl}}\varphi^+ \left(m_k \chi_k^\alpha \chi_{j\alpha} + m_l \eta^\dagger_{k\dot{\alpha}} \eta_l^{\dagger\dot{\alpha}}\right) + \text{h.c.}\ , \end{aligned} \tag{2.3}$$

where Λ_{kl} is a flavor dependent energy scale that preserves dimensions, d^c and d are SM down type quarks, and χ and η are the associated Weyl fields. The branching ratio can be obtained as

$$BR(\varphi^+ \to \bar{d^c}_k d_l) = \begin{cases} k = l & 0 \\ k \neq l & \frac{(m_k^2 + m_l^2)/\Lambda_{kl}^2}{\sum_{a=1}^{2}\sum_{b>a}(m_a^2 + m_b^2)/\Lambda_{ab}^2} \end{cases}. \tag{2.4}$$

We see that if $k = l$, the rates vanish. This can be understood in the Weyl formalism of eq. (2.3), where same flavor terms vanish due to the anti-commuting Weyl tensors and anti-symmetric color contraction.

Focusing more on the flavor structure, we choose a universal value of $\Lambda_{kl} \equiv \Lambda$ for all the down quark flavors. Hence, we find from eq. (2.4) that there is a 99.97% branching ratio to final states of φ^+ involving 1 b-quark. This is in stark contrast with the octet case, as we will soon show, where each of the scalars preferentially decays to 2 b-quarks.

Topologies with many b-quarks will already be very powerful in suppressing multijet QCD backgrounds, and an all-hadronic search could already be quite competitive. In the triplet scenario however, it will be interesting to consider the gains in sensitivity obtained by having soft leptons. This intuition will be confirmed in section 4. The only other potential source of b-jets in the triplet model come from hadronic decays of the W to b-quarks. The $W \to bt$ mode will be unaccessible because of the small scalar mass splittings we will be considering, and the off-diagonal channels will be suppressed by the CKM matrix.

2.2 Colored Octets

We now summarize the main features of the colored octet model, leaving the full details in appendix A.2. We consider a scalar doublet S transforming as $(\mathbf{8}, \mathbf{2})_{1/2}$. We can parametrize S with color index $A = 1 \dots 8$ as

$$S^A = \begin{pmatrix} S_+^A \\ \frac{S_R^A + i S_I^A}{\sqrt{2}} \end{pmatrix}. \tag{2.5}$$

Similarly to we what we did in the triplet case, we assume that S_R is lighter than $S_\pm$, and couples to down type quarks of flavor k and l via mass-weighted Yukawa couplings. With these assumptions, it is true that $m_S \gg m_k, m_l$. The branching ratios of S_R to jets is given by

$$BR(S \to \bar{d^c}_k d_l) = \begin{cases} k = l & \frac{m_k^2}{\sum_j m_j^2} \\ k \neq l & 0 \end{cases}. \tag{2.6}$$

Based on this equation, we find that decays to b-quarks will happen 99.9% of the time, and this makes the b-tags a powerful tool in the search for octet scalars.

2.3 Colored Octet-UV

Being agnostic about the UV completions of our models, it is possible that our scalars couple to colored vector-like fermions. Integrating these out generically gives rise to effective higher dimensional operators such as

$$\frac{1}{\Lambda^2} H \mathrm{Tr}\left(S G_{\mu\nu} G^{\mu\nu}\right), \tag{2.7}$$

where S is the same doublet scalar as in the octet model, H is the Higgs field, and $G_{\mu\nu}$ are gluon field strengths. This state is relevant if the interactions of eq. (2.7) dominate over those of the triplet and octet models. If this is the case, there will only be light (gluon) jets sourced by the octet-UV new physics.

3 Motivating our Choices

Throughout sections 1 and 2, we have made a number of assumptions. We have assumed a given mass hierarchy between the scalar states of the doublet, we have focused on interactions with down type quarks, and we have assumed that the best way to probe these models is via cascade decays. We now provide arguments for why each of these choices were made.

3.1 Mass Hierarchy and Interactions

Within the octet and triplet models, the scalars in the theory can a priori couple to either down type quarks, up type quarks, or a mix of the two. Since these interactions are mass weighted, interactions with up type quarks will predominantly involve top quarks, which produce very striking phenomenological signatures. Therefore the choices of model parameters that give rise to top quarks are likely to be already excluded. Another motivation to exclude top quarks in the final state is that they would produce energetic neutrinos/leptons in their semi-leptonic decays. Recall that small $\not{E}_T$ and soft leptons is one of the main features that we investigate in this paper.

We therefore choose a mass hierarchy such that the lightest scalar couples only to down type quarks. This can be achieved by picking the $-1/2$ weak eigenstate in the scalar doublet to be the lightest scalar. In the triplet model, the lightest state will be φ^+ and in the octet, it will be S_R. That said, the heavier scalar state still couples to tops. We therefore also wish to suppress the direct decay of the heavier scalars to jets, and enhance their cascade decays to the lighter scalar and a W. This can be achieved by imposing a broken Z_2 symmetry so that the heavy scalar Yukawa interactions are much more suppressed than those of the light scalar. In terms of the parameters introduced in appendix A, this translates in the requirement that $\eta_D \gg \eta_U$ and $\Lambda' \gg \Lambda$ for the octet and triplet models respectively.

3.2 Cascade Decays

Given the choices made in section 3.1, several scenarios can happen:

1. Heavy scalars form bound states instead of cascade decays, leading to dijet resonances
2. Heavy scalars cascade decay to long-lived light scalars, leading to 4 displaced tracks
3. Heavy scalars cascade decay to short-lived light scalars, leading to 4 prompt tracks

We will argue that there exists a substantial range in the parameter space of our models where scenario 3 becomes a viable option. To argue why bound states are sub-leading (scenario 1), we turn to fig. 2. In the left panel, the production cross section of pair-produced heavy scalars are shown. Also shown are upper bounds on the triplet and octet bound state cross section, which was calculated by multiplying the previous cross section by the fraction of scalar pairs that are sufficiently at rest to produce bound states. No other experimental cuts have been applied, so these bound state curves are upper bounds. Practically, using a MadGraph simulation [8], we calculated the fraction of heavy scalar pairs with total energy less than the bound state energy $|E_b|$. The bound state energy is given by [9–11]

$$|E_b| = C^2 \frac{\overline{\alpha}_s^2 m}{n^2}, \tag{3.1}$$

where C is a color factor, m is the mass of the heavy scalar, $\overline{\alpha}_s$ is the renormalized strong coupling constant and $n = 1$ for an S-wave bound state. Although clearly subleading in this plot, bound states could potentially be relevant after applying the analysis cuts that

would show up in scenario 3. To definitely settle the question, we turn to the plot in the right panel of fig. 2. There, we have applied all the analysis and signal region cuts to the triplet model for a projected sensitivity analysis that we will introduce in section 4.2. As we see, the rate for bound state production is still subleading, and so we don't consider bound states in this paper. This conclusion remains true in the case of the octet and octet-UV models, as we've explicitly checked. One caveat however is the possibility that the dijet bound state backgrounds are somehow less important, strengthening the significance of a dijet search. We have not investigated this possibility.

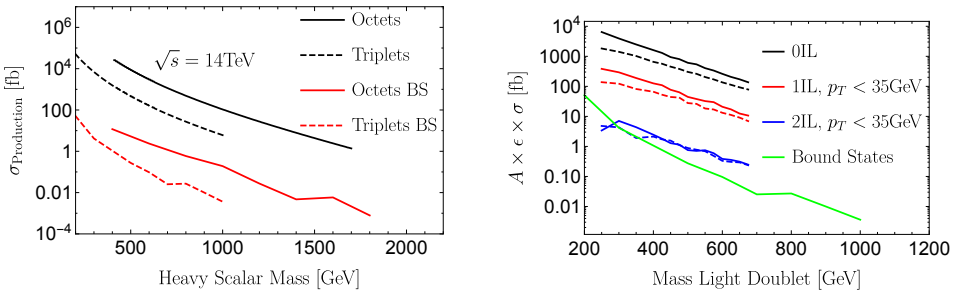

Figure 2: (Left figure) Pair production cross sections of the heavy partner colored scalars at 14TeV. For the triplet model, the underlying process is $pp \to \phi^- \phi^{-*}$. For the octet model, the underlying process is $pp \to s^+ s^-$. The bound states curves (BS) denote the same respective processes accounting for the efficiency of producing the two heavy scalars with total energy less than their bound state energy.
(Right figure) Triplet model comparison between the bound state cross section and the $A \times \epsilon \times \sigma$ for the projected analysis described in section 4.2. Solid (dotted) lines indicate a 10GeV (70GeV) mass splitting. In all cases, the bound state rates are subdominant.

We now turn our attention to the possibility of obtaining displaced vertices (scenario 2). Taking the octet model for the purposes of this discussion, the parameters η_u and η_d control both the decay length of the scalars as well as the branching ratio of the heavy scalar state into an (off-shell) W boson and the light scalar state. These dependences are calculated explicitly in fig. 3. In order to have the best possible prospects for detecting scalar doublets via scenario 3, we would like to maximize the branching ratio (shown in the left panel), while keeping the $c\tau$ (shown in the right panel) below the 1mm level. This can be done for values of η as low as 10^{-7}, and as high as 10^{-4} or 10^{-1} for mass splittings of 10GeV or 80GeV respectively. For the remainder of the paper, we will assume that the values of η are within these ranges, and that the cascade decay branching ratios to leptons and jets saturate to those of the W boson. We note that with these choices, the production of the heavy scalars proceeds via the SU(3) strong coupling constant, and the cascade decays proceed via the SU(2) weak coupling constant, with no free parameters to be determined. Having motivated a search based on cascade decays, we now proceed to developing a search strategy.

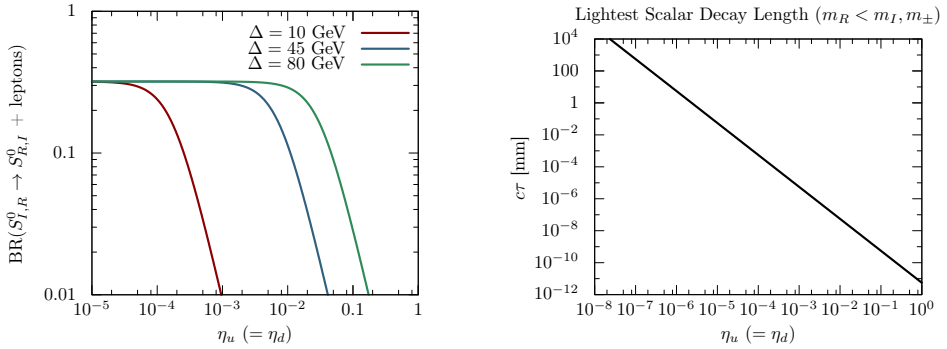

Figure 3: Branching ratio (left) and decay length (right) versus the coupling $\eta_{u,d}$ for the Octet model. Branching ratios are shown for several characteristic mass splittings Δ between the heavy and light partners of the scalar doublet. There is a significant range in $\eta_{u,d}$ for which the leptonic branching ratio of the heavy scalar is equal to that of the W boson, while ensuring the light scalar partners decay promptly.

4 New Signal Regions

We previously presented several models of colored scalar doublets and motivated a region in parameter space where we can search for them based on cascade decays. Our main interest now lies in the very interesting possibility that the colored doublets exhibit a compressed mass spectrum, with a splitting smaller than the W mass. This leads to the novel possibility of producing soft leptons from off-shell W bosons, many jets and small $\not{E}_T$. Such a topology is hidden from conventional searches, as we will confirm, and so additional signal regions not often presented in LHC searches are warranted. These signal regions are motivated in our projected sensitivity analysis in section 4.2. In section 4.3, we recast a conventional search without soft lepton/low $\not{E}_T$ cuts. As we will see, substantially better sensitivity to scalar doublets in the compressed mass regime can be achieved in the former case. We will also explain in what sense these results are robust to how the hadronic content of the events are treated.

The dominant backgrounds associated to this topology depend on the requirements made on the number of b-tags and the number of leptons. A full detailed list for each channel in our projected sensitivity analysis will be provided in table 4. Generally speaking, the dominant backgrounds will be either $t\bar{t}$ + jets (when requiring b-tags with or without leptons), QCD (when requiring no leptons) and Z/W + jets (when requiring 1 or more leptons). We will show that for the backgrounds, the lepton p_T and event $\not{E}_T$ follow broad kinematic distributions. For our signals of interest, they will almost always favor very small values and we will try to capitalize on these differences.

There is another potential source of backgrounds arising from objects that fake low p_T leptons. These fake backgrounds would need to be evaluated by the LHC collaborations,

as we don't have the tools to estimate the size of their contribution. To limit their potential impact, we will require that the lepton p_T is larger than 8GeV, and that the leptons pass both a relative I_{rel} and absolute I_{abs} isolation cut. There is another potential way to limit this background if it is still significant after these requirements. In signal events, low lepton p_T and low event $\not{E}_T$ are highly correlated. For these fake backgrounds, if it turns out to be the case that events that give rise to fake leptons also have large $\not{E}_T$, then the $\not{E}_T$ cuts we propose could be significantly tightened with little effect on the signal rates.

All of our simulations at parton level are done with MadGraph v2.4.0 and v2.6.1 [8]. For the projected sensitivity analysis, we use the "nn23lo1" parton distribution function and for the recast analysis, we use "NNPDF30_nlo_as_0118". We have implemented our signal models using FeynRules 2.3 [12]. Parton level events are showered and hadronized using Pythia 8212 [13]. Matching, when performed for background events, is done with the shower-k_T algorithm. Event analysis is performed using our own HEPMC-based implementation. Jets are clustered using the anti-k_T algorithm [14] with a 0.5 and 0.4 jet radius parameter R, for the projected and recast analyses respectively.

4.1 Baseline Hadronic Analysis

In order to study the gains in sensitivity achieved by requiring soft isolated leptons and low $\not{E}_T$, we made a "signal region study" (SR Study). In this study, we simulated at 14TeV all of the signals and backgrounds for the 0, 1 and 2 lepton channels of our models. The hadronic component of the analysis is motivated largely from an 8TeV hadronic CMS search [1], which we refer to as the baseline analysis. These are shown in table 1. On top of this baseline analysis, we add leptonic and $\not{E}_T$ requirements, which we refer to as our signal regions.

Observable	Values
$\sqrt{s}$	14TeV
Fourth jet p_T	>100GeV
$\Delta m/m_{\text{avg}}$	<0.15
$\overline{\Delta}$	>100GeV
$\Delta\eta_{\text{dijet}}$	<1.0
N_{Jets}	≥4
I_{rel}	< 0.5
I_{abs}	< 5GeV

Table 1: Baseline cuts and lepton isolation requirements for the *SR study*. These cuts are in part motivated from the 8TeV CMS search [1], which we re-simulated at 14TeV for 35.9fb^{-1} of data.

We now explain how the baseline analysis is made. Collecting events with 4 or more jets, we identify pairs of jets by minimizing the measure

$$\Delta R_{\text{dijet}} = \sum_{i=1,2} |\Delta R^i - 1|. \tag{4.1}$$

In signal events, these tend to identify jets that originate from the same parent resonance. The variables Δm and $m_{\rm avg}$ can then be constructed by respectively taking the difference and average of the masses of the two pairs of dijets. Finally, the variable $\overline{\Delta}$ is constructed as

$$\overline{\Delta} = \left(\sum_{i=1,2} |p_T^i| \right) - m_{\rm avg}, \tag{4.2}$$

where the sum is over the transverse momentum of the two dijet pairs. We calibrated our simulations and baseline analysis by reproducing the "Inclusive data (high-mass search)" plot from Figure 5 of [1] at 8TeV.

We now discuss leptons. In order to minimize the effect of background contamination from fake leptons, we imposed lepton isolation cuts. Our lepton isolation criteria is based on those of [15]. For a lepton ℓ, we define the relative lepton isolation as

$$I_{\rm rel}(\ell) = \sum_{\substack{i \neq P \\ p_T > 1{\rm GeV} \\ \Delta R < 0.4}} \frac{p_T(i)}{p_T(\ell)}, \tag{4.3}$$

where the sum is over all particles i with transverse momentum larger than 1GeV and within a $\Delta R < 0.4$ radius. We define the absolute isolation as $I_{\rm abs} = p_T(\ell) I_{\rm rel}(\ell)$. The cuts on these variables are shown in table 1. We also impose detection efficiencies on the leptons as a function of p_T and η as found in the Delphes 3.4.1 CMS configuration card [16].

As discussed in section 2, the number of b-tags we imposed in our sensitivity estimates depend on the model. In general, the triplets will tend to produce 2 b-quarks and the octets will tend to produce 4 b-quarks. As for the octet-UV extension, we only expect the new particles in the theory to produce light jets. The b-tag efficiencies we have assumed are parametrized as [16]

$$\epsilon = 0.85 \tanh(0.0025 p_T) \left(\frac{25.0}{1 + 0.063 p_T} \right). \tag{4.4}$$

The various requirements on the number of b-tags and number of leptons will determine which backgrounds are dominant for a given signal. In table 4, we show the b-tags, signals and backgrounds that contribute in each of the channels we have simulated.

4.2 Signal Regions

Now that we have a baseline analysis and a handle on leptons, we can determine the signal regions by studying differences in the signal and background kinematics. Collecting events with one or more isolated leptons, we plot normalized momentum distributions of the highest-p_T lepton and of the event $\not{E}_T$. This is done in fig. 4 for the triplet model, and in fig. 7 for the octet and octet-UV models. In all cases, we notice that when the scalar doublet mass splitting is smaller than the mass of the W boson, the signal kinematics differ largely from those of the dominant backgrounds. This is mainly due to the fact that in the

$t\bar{t}$ + jets and W + jets backgrounds, the W tends to be produced on-shell and gives rise to a much harder spectrum, as compared to the signal. Based on these kinematic results, we propose to add signal region cuts on the lepton p_T and the $\not{E}_T$ of the event. These are shown in table 2 for the various leptonic channels.

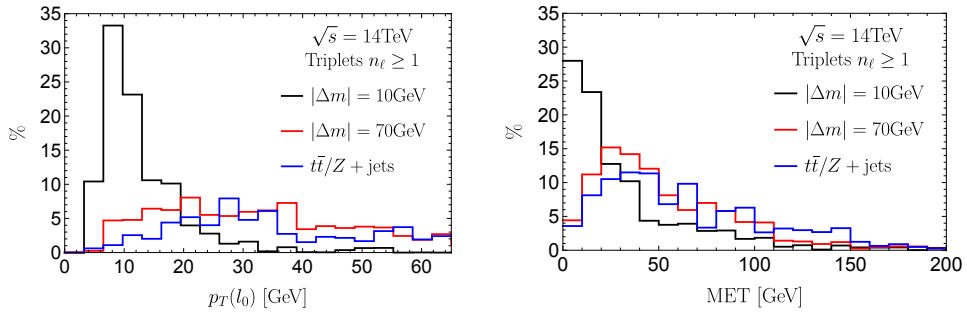

Figure 4: Normalized kinematic distributions of leading lepton p_T (left) and missing energy (right) for a 500GeV light partner *triplet* scalar and combined $t\bar{t}$/Z + jets background. The distributions shown are for events containing at least 1 lepton and which satisfy all of the *SR study* selection requirements (besides the missing energy and lepton p_T cuts). Equivalent plots for the octet and octet-UV models are shown in fig. 7.

The signal region cuts proposed in table 2 and the lepton isolation criteria will naturally lower our signal statistics as compared with conventional searches. Later on in our sensitivity plots, we will show that this additional penalty is acceptable given the much larger penalty that will be incurred by the background. It is interesting nevertheless to quantify the fraction (which we will refer to as ϵ) by which the signal statistics drop given the signal region cuts. Since this fraction is similar between all our models and depend primarily on the doublet mass splitting and not the overall masses of the scalars (as we've numerically checked), we discuss only the triplet model with a 10GeV (70GeV) mass splitting. Starting from the topology in fig. 1 and allowing all possible decays of the W, the ϵ after requiring the absence of isolated leptons is 90% (82%). The larger ϵ for a smaller mass splitting is reasonable, since soft leptons are more likely to get lost and not pass the isolation cuts. Next we consider events with at least 1 lepton that pass the baseline cuts. From those events, requiring at least 1 lepton to be isolated with $p_T \geq 8$GeV gives an ϵ of 21% (49%). Imposing in addition the *1 iso. lepton* signal region cuts in table 2 brings ϵ to 19% (25%). As we see, signals with a small mass splitting are not affected by the signal region cuts. The signals with larger splitting, and by extension the backgrounds, are very sensitive to this cut. For events with 2 leptons, requiring that these both be isolated with $p_T \geq 8$GeV gives an ϵ of 4% (22%). Imposing the additional *2 iso. lepton* cuts brings ϵ to 3% (7%). Although ϵ for small mass splittings changes only marginally after the *2 iso. lepton* cuts have been imposed, the overall statistics in the 2 lepton channel are becoming very low as compared to the 0 or 1 lepton channels, and we might expect the sensitivity to suffer as a result.

0 iso. lepton	**1 iso. lepton**	**2 iso. leptons**
No isolated leptons	$p_T(\ell^\pm) \leq 40\text{GeV}$	$p_T\{\ell^+,\ell^-\} \leq \{40,30\}\text{GeV}$
	$\not{E}_T \leq 90\text{GeV}$	$\not{E}_T \leq 90\text{GeV}$
	$p_T(\ell^\pm) \geq 8\text{GeV}$	$p_T(\ell^\pm) \geq 8\text{GeV}$
		$\lvert m_{\ell^-+\ell^+} - m_Z\rvert \geq 10\text{GeV}$

Table 2: Signal region cuts proposed in this paper.

4.3 Recast Study

Having developed several signal regions based on the presence of soft isolated leptons, we would like to see how they compare to a conventional search. The 1 lepton channel seemed particularly powerful in the small mass splitting limit, providing both acceptable signal efficiency and strong background rejection. To provide a benchmark comparison for the cuts we are advocating, we will recast an existing search with a topology of many jets and 1 lepton. Many such searches that have already been performed require large $\not{E}_T$, motivated in part by the presence of a light stable particle that escapes the detector. Looking at our characteristic signal kinematics in fig. 4, a large $\not{E}_T$ requirement will preclude any sensitivity to our model. Some searches in the literature however have low or modest $\not{E}_T$ requirements [2, 17–21]. Of these, we will recast the 13TeV CMS search [21]. For convenience, we show the main features of their analysis in table 3, which we implemented exactly for our triplet and octet models. In order to calibrate our analysis code, we simulated a $t\bar{t}+\text{jets}$ matched sample and reproduced to within 30% the post-analysis rates shown in Table 1 of [21].

Analysis Cuts	
$\sqrt{s}$	13TeV
$\mathcal{L}$	35.9fb^{-1}
# leptons	1 $e^\pm$ or 1 $\mu^\pm$
$p_T(\ell)$	$> 20\text{GeV}$
$\lvert\eta(\ell)\rvert$	$< 2.5(e^\pm)$, $< 2.4(\mu^\pm)$
$I_{\text{rel}}(\ell)$	$< 0.1(e^\pm)$, $< 0.2(\mu^\pm)$
jets	$p_T > 30\text{GeV}$, $\lvert\eta\rvert \leq 2.4$
# jets, # b-tags	≥ 4, ≥ 1
$\not{E}_T$	$\geq 0\text{GeV}$
H_T	$> 1200\text{GeV}$
M_J	$> 500\text{GeV}$

Table 3: *Recast study* analysis cuts from the 13TeV CMS search [21].

After simulating our signal according to the cuts in table 3, we can study the normalized lepton p_T and $\not{E}_T$ kinematic distributions in events containing exactly 1 lepton. These are shown in fig. 5 for a 400GeV light partner triplet scalar, and in fig. 8 for a 1TeV light partner octet scalar. In both cases, we see a very large peak above the dominant background for signal events with a small mass splitting. This indicates that in the compressed mass

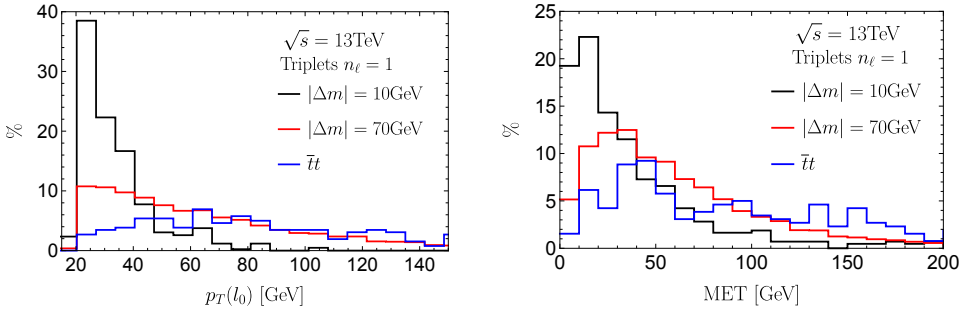

Figure 5: Normalized kinematic distributions of leading lepton p_T (left) and missing energy (right) for a 400GeV light partner *triplet* scalar and $t\bar{t}+$jets background. The distributions shown are for events containing exactly 1 lepton and which satisfy all of the *recast study* selection requirements.

regime, conventional analyses are not taking full advantage of the differences in signal and background kinematics.

5 Results

To summarize, we have so far considered 3 different colored scalar doublet models which have the capacity to exhibit a compressed mass spectrum. With this interesting limit in mind, we re-simulated at 14TeV a hadronic baseline analysis [1], on top of which we investigated features of the 0, 1 and 2 isolated lepton channels. Stark kinematic differences between signals and backgrounds motivated us to apply additional lepton p_T and $\not{E}_T$ signal region cuts shown in table 2, with very little effect on the signal rates in the compressed limit. Due to surprisingly good signal efficiencies in the 1 isolated lepton channel, we recasted a new search [21] that looked for 1 lepton and many jets, with no $\not{E}_T$ requirements or additional lepton p_T cuts.

We are now in a position to show the main results of this paper. In fig. 6, we calculate the gains in sensitivity across each of the lepton channels and signal regions, and how they compare to the analysis we recasted. The sensitivity results shown here account for everything; namely the luminosity, production cross section, branching ratio, as well as the efficiencies from the object selection, event analysis and signal regions. We choose to calculate the sensitivity according to the procedure in appendix B, since it has the properly that it doesn't diverge in the low background limit as compared to usual $s/\sqrt{b}$ estimates. We now discuss the sensitivity plots in fig. 6. There are 2 main features we wish to emphasize. The first is the interplay between the number of b-tags and the power of the 0 and 1 lepton channels in the projected signal significance plots (left column). As a reminder of the information in table 4, the triplet plot requires 2 b-tags and the octet-UV plot requires 0 b-tags. For the octet model, the 0 lepton channel requires 4 b-tags and the other ones require 2. The larger the number of b-tags, the more aggressively we cut down on the QCD

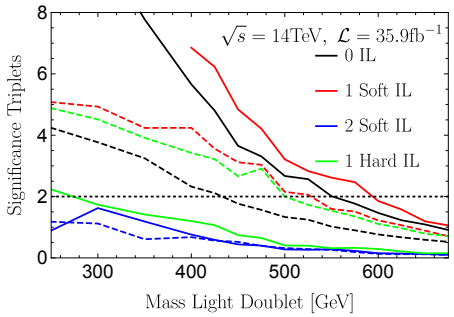

(a) SR Study, Triplet

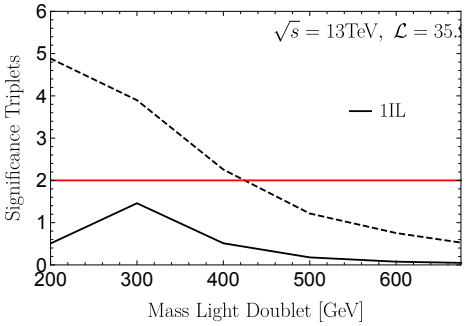

(b) Recast Study, Triplet

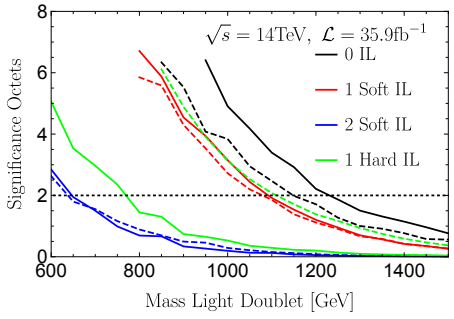

(c) SR Study, Octet

In preparation

(d) Recast Study, Octet

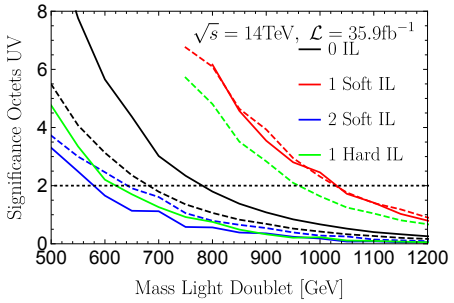

(e) SR Study, Octet UV

In preparation

(f) Recast Study, Octet UV

Figure 6: Projected signal significance (left column) and recasted exclusions of an existing search [21] (right column) with 35.9fb^{-1} of data at the LHC. Results are shown for the triplet model (top row), octet model (middle row) and octet-UV model (bottom row). Soft IL (isolated leptons) denote that the cuts in table 2 have been applied, and Hard IL denotes that the lepton is required to have $p_T > 20$GeV. Solid (dotted) lines denote mass splittings of 10GeV (70GeV) between the heavy and light partners of the scalar doublet. All event selection cuts are described in section 4. When searching for theories featuring cascade decays with a compressed mass spectrum, these sensitivity plots demonstrate the power of adding simple signal regions that target events with soft leptons and low missing energy.

background and as we see, this can give better sensitivity than requiring isolated leptons. When the flavor content of our signal is not suited to a large number of b-tagged states, requiring 1 isolated lepton can provide an improvement over the all-hadronic scenario. For the octet-UV model, which contains only light jets, this improvement is particularly accentuated.

The second feature we wish to discuss is the weakness of the 2 lepton channel. Requiring just the presence of 2 isolated leptons imposes a large penalty on signals with small mass splitting (and in signals with large mass splitting after requiring soft leptons). Because of this, the signal rates drop quickly as a function of mass below 1 event (assuming 35.9fb^{-1}), and the sensitivity doesn't benefit from the even larger background reduction. The two lepton channel could still be of interest in cases where there are unusually large systematic uncertainties associated to having too many jets in the final state, as a tool for cross-validation in light of evidence for new physics, or in final states such as those of the octet-UV model with few b-tagged states.

We now discuss the results of the recast (right column), where we show data driven constraints obtained from our recast. To make the connection with our projected sensitivity analysis, we have calculated a "hard lepton curve" (green curve in the left column) in the 1 isolated lepton channel that only requires the lepton to have $p_T > 20\text{GeV}$ and makes no mention of the $\not{E}_T$ in the event. This is very similar to how the lepton is treated in the recasted analysis and serves to compare the two analyses. As we see, the projected sensitivities achieved with this hard lepton channel are very similar to the sensitivities from the recast, despite the very different methods of treating the hadronic content of the event. Furthermore, the kinematics of the recast analysis, shown in figs. 5 and 8, also suggests important differences between the signals and backgrounds that are not being capitalized. These observations demonstrate the robustness of the soft lepton/low $\not{E}_T$ sensitivity gains to changes in the hadronic part of the analysis.

6 Discussion and Conclusions

In conclusion, we have motivated a new search topology at the LHC by considering several new models and demonstrating that their limits are subject to improvement when requiring soft leptons and low missing energy. In order for this conclusion to hold now, we would need to redo our analysis using the updated search for pair produced resonances [22] at 13 TeV. Although this search has a lot in common with the 8 TeV search [1], it does feature more stringent cuts and also introduces new variables on which cuts are imposed. It will interesting to know if the kinematic differences between signal and background in the missing energy and lepton p_T distributions disappear under this new analysis. If they do not, then this topology should also further reduce backgrounds and provide improvement in sensitivity. We are currently investigating this question.

A Full Model Details

A.1 Colored Triplets

We consider a scalar Φ in the $(\mathbf{3}, \bar{\mathbf{2}})_Y$ representation. We denote the doublet components generically as $\Phi = (\varphi^-, \varphi^+)$, although the specific electric charges depend on the value of Y. The scalar is produced through its strong coupling. It can decay only via coupling to quarks. Φ must couple to Q to absorb the doublet transformation. ΦQ transforms as a $\bar{\mathbf{3}}$ under $\mathrm{SU}(3)_\mathrm{c}$ and cannot couple to any other Weyl spinor. Therefore, the lowest-dimension all-hadronic operators are

$$\mathcal{L} \supset \frac{1}{\Lambda_{ij}} \Phi (d_i^\mathrm{c})^\dagger \bar{\sigma}^\mu \overleftrightarrow{D_\mu} Q_j + \frac{1}{\Lambda'_{ij}} \Phi' (u_i^\mathrm{c})^\dagger \bar{\sigma}^\mu \overleftrightarrow{D_\mu} Q_j + \mathrm{h.c.}, \tag{A.1}$$

for fields with $Y(\Phi) = +1/6$ and $Y(\Phi') = -5/6$. The gauge field part of the covariant derivative cancels when taking the difference of the left and right action D^μ, so we can replace $D^\mu \to \partial^\mu$. The color indices are implicitly contracted with an antisymmetric tensor. Note that, since the Weyl spinors are different, the flavor indices i and j can be the same. In terms of the corresponding Dirac spinors,

$$\mathcal{L} \supset \frac{1}{\Lambda_{ij}} \Phi \bar{d}_i^\mathrm{c} \overleftrightarrow{\partial\!\!\!/} \mathcal{P}_\mathrm{L} Q_j + \frac{1}{\Lambda'_{ij}} \Phi' \bar{u}_i^\mathrm{c} \overleftrightarrow{\partial\!\!\!/} \mathcal{P}_\mathrm{L} Q_j + \mathrm{h.c.} \tag{A.2}$$

We can use the equations of motion,

$$\partial_\mu \bar{\psi} \gamma^\mu = i M_\psi \bar{\psi} - i g \bar{\psi} A\!\!\!/^a T^a + \ldots, \tag{A.3}$$

$$\partial\!\!\!/ \psi = -i M_\psi \psi + i g A\!\!\!/^a T^a \psi + \ldots, \tag{A.4}$$

where $\ldots$ represents the contribution of the ψ Yukawa couplings. Because the mass term leads to two-body Φ decay and all the other terms give three-body decay, we only keep the leading mass term. Using the equations of motion and $\{\gamma^\mu, \gamma^5\} = 0$, we get (in $\mathrm{SU}(2)_\mathrm{L}$

components)

$$\begin{aligned}\mathcal{L} \supset & -\frac{i}{2\Lambda_{ij}}\varphi^{-}\bar{d}_i^{\rm c}\left[(M_{d_i}+M_{u_j})-\gamma^5(M_{d_i}-M_{u_j})\right]u_j \\ & -\frac{i}{2\Lambda_{ij}}\varphi^{+}\bar{d}_i^{\rm c}\left[(M_{d_i}+M_{d_j})-\gamma^5(M_{d_i}-M_{d_j})\right]d_j \\ & -\frac{i}{2\Lambda'_{ij}}(\varphi')^{-}\bar{u}_i^{\rm c}\left[(M_{u_i}+M_{u_j})-\gamma^5(M_{u_i}-M_{u_j})\right]u_j \\ & -\frac{i}{2\Lambda'_{ij}}(\varphi')^{+}\bar{u}_i^{\rm c}\left[(M_{u_i}+M_{d_j})-\gamma^5(M_{u_i}-M_{d_j})\right]d_j \\ & +\text{h.c.}\end{aligned} \tag{A.5}$$

We see that the couplings are mass-proportional; thus, we expect operators involving top quarks to dominate. We also should consider the flavor structure of Λ_{ij}. Using the conventions for the flavor group $\mathrm{SU}(3)_Q \times \mathrm{SU}(3)_d$ where $d^{\rm c}$ transforms as $(\mathbf{1},\mathbf{3})$ and Q transforms as $(\bar{\mathbf{3}},\mathbf{1})$, the spurion Y_d transforms as $(\mathbf{3},\bar{\mathbf{3}})$, since $Y_d Q d^{\rm c}$ is invariant. In our operator, we have $(d^{\rm c})^\dagger Q$, which transforms as $(\bar{\mathbf{3}},\bar{\mathbf{3}})$. There is no way of inserting Y_d to absorb the flavor rotation, so this operator is forbidden in minimal flavor violation. Of course, if Λ is of sufficiently high scale, this is not a problem, but it means that there is not necessarily a preferred flavor structure. The most generic thing to consider is a flavor-universal configuration, which implicitly gives dominant coupling to t and b because of the mass insertions from the Dirac equation. The same arguments trivially follow for the coupling Λ'_{ij}.

A.2 Colored Octets

We begin by writing the Yukawa and Kinetic sectors of the Lagrangian of a $(\mathbf{8},\mathbf{2})_{1/2}$ scalar field:

$$\begin{aligned}\mathcal{L} =& \mathrm{Tr}[(D_\mu S)^\dagger (D_\mu S)] - \eta_D g^D_{ij}\bar{d}_{Rj}T^A(S^A)^\dagger Q_i - \eta_U g^U_{ij}\bar{u}_{Rj}T^A(i\sigma_2 S^{*A})^\dagger Q_i + \text{h.c.} \\ & + 2m_S^2\mathrm{Tr}S^{\dagger i}S_i + \lambda_1 H^{\dagger i}H_i\mathrm{Tr}(S^{\dagger j}S_j) + \lambda_2 H^{\dagger i}H_j\mathrm{Tr}(S^{\dagger j}S_i) \\ & + [\lambda_3 H^{\dagger i}H^{\dagger j}\mathrm{Tr}(S_iS_j) + \lambda_4 H^{\dagger i}\mathrm{Tr}(S^{\dagger j}S_jS_i) + \lambda_5 H^{\dagger i}\mathrm{Tr}(S^{\dagger j}S_iS_j)] + \text{h.c.} \\ & + \lambda_6\mathrm{Tr}(S^{\dagger i}S_iS^{\dagger j}S_j) + \lambda_7\mathrm{Tr}(S^{\dagger i}S_jS^{\dagger j}S_i) + \lambda_8\mathrm{Tr}(S^{\dagger i}S_i)\mathrm{Tr}(S^{\dagger j}S_j) \\ & + \lambda_9\mathrm{Tr}(S^{\dagger i}S_j)\mathrm{Tr}(S^{\dagger j}S_i) + \lambda_{10}\mathrm{Tr}(S_iS_j)\mathrm{Tr}(S^{\dagger i}S^{\dagger j}) + \lambda_{11}\mathrm{Tr}(S_iS_jS^{\dagger j}S^{\dagger i}).\end{aligned} \tag{A.6}$$

Writing out the covariant derivative gives:

$$\begin{aligned}D_\mu S &= \partial_\mu S - ig_W W^a_\mu \tau^a S - i\frac{g_Y}{2}B_\mu S + g_s i[G_\mu, S] \\ &= \left(\partial_\mu S^A - ig_W W^a_\mu \tau^a S^A - i\frac{g_Y}{2}B_\mu S^A - g_s f_{ABC}G^B_\mu S^C\right)T^A.\end{aligned} \tag{A.7}$$

To explain the notation, we have $S = S^A T^A$ where S is a doublet under SU(2), $S^i = S^{iA}T^A$ (for the scalar potential terms) where i is the SU(2) index, and $G_\mu = G^A_\mu T^A$. T^A are the 8 Gell-Mann matrices whose matrix components run from 1 to 3, and $\eta_{D,U}$ are complex proportionality constants that arise when requiring that the Yukawas become diagonal in

the quark mass basis. The kinetic term in eq. (A.6) will have an overall coefficient of the form $\mathrm{Tr}[T^A T^D] = \frac{\delta^{AD}}{2}$, where $1/2$ is the index of the fundamental representation.

Next, we express

$$S^A = \begin{pmatrix} S_+^A \\ \frac{S_R^A + i S_I^A}{\sqrt{2}} \end{pmatrix}. \tag{A.8}$$

We also rotate the flavor of the quarks to the mass basis, and express the Yukawa couplings in terms of the quark masses. After all these transformations, the Lagrangian we obtain is:

$$\begin{aligned} \mathcal{L} \;=\; & \mathrm{Tr}[(D_\mu S)^\dagger (D_\mu S)] \\ & - \eta_D \frac{\sqrt{2} m_{ij}^d}{v} \bar{d}_{Ri} V_{jk}^\dagger T^A S_-^A u_k - \eta_D \frac{m_{ij}^d}{v} \bar{d}_{Ri} T^A (S_R^A - i S_I^A) d_j \\ & - \eta_U \frac{m_{ij}^u}{v} \bar{u}_{Ri} T^A (S_R^A + i S_I^A) u_j + \eta_U \frac{\sqrt{2} m_{ij}^u}{v} \bar{u}_{Ri} V_{jk} T^A S_+^A d_k \\ & + \text{h.c.} + \text{Quadratic} + \text{Quartic} + \text{Higgs}. \end{aligned} \tag{A.9}$$

Here, m_{ij}^u and m_{ij}^d are the diagonal up and down type quark mass matrices respectively, and V is the CKM matrix. We might expect mixing between S_R^A and S_I^A. From eq. (A.6), the only terms that could lead to mixing are the terms proportional to m_S^2, λ_1, λ_2, and λ_3. After spontaneous symmetry breaking of the Higgs, we get no mixing between S_R^A and S_I^A. We find that the masses are

$$\begin{aligned} m_{S^\pm}^2 &= m_S^2 + \frac{\lambda_1 v^2}{4} \\ m_{S_R}^2 &= m_S^2 + \frac{v^2}{4}(\lambda_1 + \lambda_2 + 2\lambda_3) \\ m_{S_I}^2 &= m_S^2 + \frac{v^2}{4}(\lambda_1 + \lambda_2 - 2\lambda_3). \end{aligned} \tag{A.10}$$

B Sensitivity Calculation

We briefly outline how the sensitivities in our plots were calculated. The strategy is based on the 2009 PDG on statistics [23]. We consider a counting experiment where the experiment has seen n events, whereas b were predicted from the Standard Model and s from new physics (for projected sensitivities we assume that $n = b$). In a Bayesian framework given a posterior probability p and likelihood function L, one can define a $(1-\alpha)$ credibility level as

$$1 - \alpha = \int_0^s p(s|n) ds = \frac{\int_{-\infty}^s L(n|s)\pi(s) ds}{\int_{-\infty}^\infty L(n|s)\pi(s) ds}. \tag{B.1}$$

Using a flat prior π in the new physics signal rate and the Poisson likelihood function

$$L(n|s) = \frac{(s+b)^n}{n!} \mathrm{e}^{-(s+b)}, \tag{B.2}$$

eq. (B.1) can be rewritten as

$$\begin{aligned}\alpha &= \mathrm{e}^{-s}\frac{\sum_{m=0}^{n}(s+b)^m/m!}{\sum_{m=0}^{n}b^m/m!} \\ &= \frac{\Gamma_{\text{upper incomplete}}(1+n, b+s)}{\Gamma_{\text{upper incomplete}}(1+n, b)}.\end{aligned} \tag{B.3}$$

The number of signal events s is calculated as

$$s = \mathcal{L}\sigma\epsilon A, \tag{B.4}$$

where $\mathcal{L}$ is the luminosity of the experiment, σ denotes the total cross section for the processes of interest (including branching ratios), ϵ denotes the efficiency of the analysis cuts, and A denotes the acceptance of the final states within the detector. It is important to note that in our setup, σ is completely determined in terms of the electroweak coupling and the branching of the W bosons, with no overall new coupling constant. Calculating $\alpha \in [0, 1]$ indicates how significant the signal is over the background, with values of α below 5% considered as significant, and hence excluded. In all our sensitivity plots, we show the quantity a instead of α, by doing a $1-1$ mapping to a Gaussian distribution. This is done in order to make the results more intuitive; we are not implying that the data should follow a Gaussian distribution. Explicitly, a is obtained as a function of α by solving the injective equation

$$1 - \alpha = \frac{1}{\sqrt{2\pi}}\int_{-a}^{a}\exp\left(-\frac{x^2}{2}\right)\mathrm{d}x, \tag{B.5}$$

where we have normalized the Gaussian average and standard deviation to 0 and 1 respectively. The exclusion criteria of $\alpha < 5\%$ now corresponds to $a > 1.96$.

Using these methods is desirable over the simpler $s/\sqrt{b}$ and s/b metrics, which overestimate the sensitivity in the limit of 0 background. Interestingly, setting $(1-\alpha) = 95\%$CL in eq. (B.1) for a 0 background search implies that the criteria for exclusion is the commonly cited $s = 3$ signal events.

C Additional Details

C.1 SR Study

Processes	0 iso. lepton	1 iso. lepton	2 iso. leptons
Triplets Search			
B-tags	$n_{\text{b-tags}} \geq 2$	$n_{\text{b-tags}} \geq 2$	$n_{\text{b-tags}} \geq 2$
Signal $\phi^-\phi^{-*} \to$	$(4j)(2j\ell^\pm)(\ell^\mp\ell^\pm)\phi^+_{j\bar{b}}\phi^{+*}_{jb}$	$(2j\ell^\pm)(\ell^\mp\ell^\pm)\phi^+_{j\bar{b}}\phi^{+*}_{jb}$	$\ell^+\ell^-\phi^+_{j\bar{b}}\phi^{+*}_{jb}$
Backgrounds	$t\bar{t}+0,1,2J$ $2b+1,2J$	$t\bar{t}+0,1,2J$	$t\bar{t}+1,2J$ $L^+L^-2b+1,2J$
Octets Search			
B-tags	$n_{\text{b-tags}} \geq 4$	$n_{\text{b-tags}} \geq 2$	$n_{\text{b-tags}} \geq 2$
Signal $s^+s^- \to$	$(4j)(2j\ell^\pm)(\ell^\mp\ell^\pm)s_{\bar{b}b}\bar{s}_{\bar{b}b}$	$(2J\ell^\pm)(\ell^\mp\ell^\pm)s_{JJ}\bar{s}_{JJ}$	$\ell^+\ell^-s_{JJ}\bar{s}_{JJ}$
Backgrounds	$t\bar{t}+2b$ $4b$	$t\bar{t}+0,1,2J$	$t\bar{t}+1,2J$ $L^+L^-2b+1,2J$
Octets UV Search			
B-tags	$n_{\text{b-tags}} \geq 0$	$n_{\text{b-tags}} \geq 0$	$n_{\text{b-tags}} = 0$
Signal $s^+s^- \to$	$(4J)(2J\ell^\pm)(\ell^\mp\ell^\pm)s_{gg}s_{gg}$	$(2J\ell^\pm)(\ell^\mp\ell^\pm)s_{gg}s_{gg}$	$\ell^+\ell^-s_{gg}s_{gg}$
Backgrounds	$W^\pm+3,4J$ $3,4J$	$W^\pm+3,4J$	$L^+L^-+3,4j$

Table 4: Signals and backgrounds we have simulated at $\sqrt{s} = 14$TeV in our SR study. For each search and signal region, we indicate requirements on the number of b-tags, missing energy and lepton p_T. The symbol j denotes all light quarks (J includes bottom quarks), ℓ denotes light leptons (L includes taus), and subscripts indicate specific final state decays. A sequence of states in brackets indicate several independent contributions to a given signal. Neutrinos are not written and particles shown are not assumed to be on-shell.

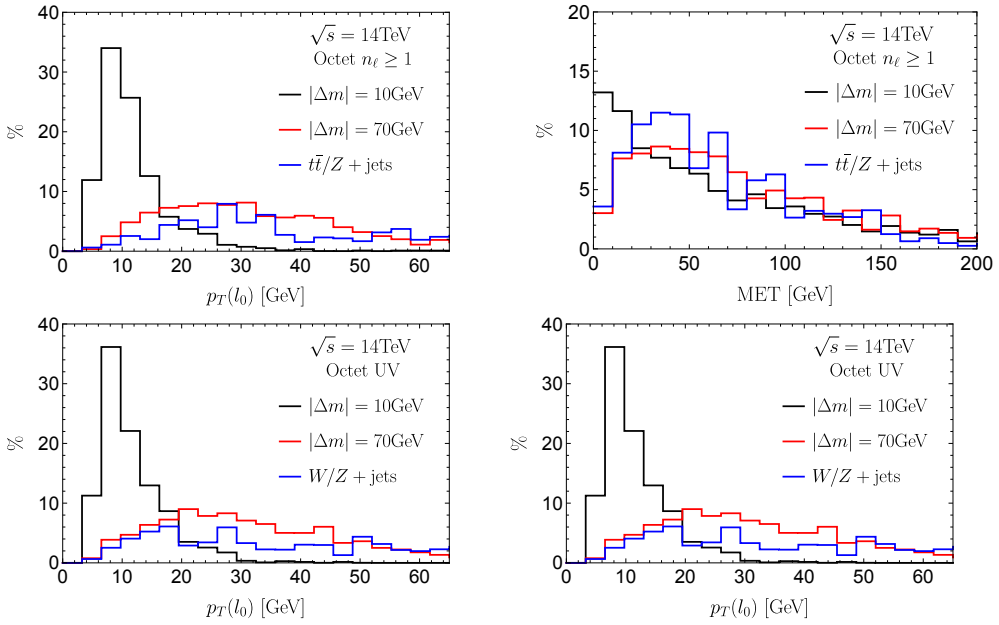

Figure 7: Normalized background and signal kinematic distributions of leading lepton p_T (left column) and missing energy (right column). We show results for a 1TeV light partner *octet* model (top row) and a 750GeV light partner *octet-UV* model (bottom row). The distributions shown are for events containing at least 1 lepton and which satisfy all of the *SR study* selection requirements (besides the missing energy and lepton p_T cuts).

C.2 Recast Study

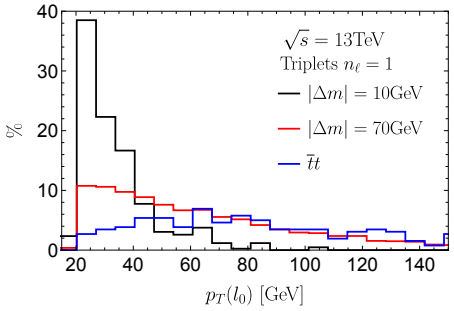

In preparation

Figure 8: Normalized kinematic distributions of leading lepton p_T (left) and missing energy (right) for a 1TeV light partner *octet* scalar and $t\bar{t}$ + jets background. The distributions shown are for events containing exactly 1 lepton and which satisfy all of the *recast study* selection requirements.

Acknowledgments

We thank Itay Yavin and Cliff Burgess for useful discussions. This research was supported in part by Perimeter Institute for Theoretical Physics. Research at Perimeter Institute is supported by the Government of Canada through the Department of Innovation, Science and Economic Development and by the Province of Ontario through the Ministry of Research and Innovation. This research was also supported by funds from the National Science and Engineering Research Council of Canada (NSERC).

References

[1] Vardan Khachatryan et al. Search for pair-produced resonances decaying to jet pairs in proton–proton collisions at $\sqrt{s} = 8$ TeV. *Phys. Lett.*, B747:98–119, 2015.

[2] Vardan Khachatryan et al. Search for stealth supersymmetry in events with jets, either photons or leptons, and low missing transverse momentum in pp collisions at 8 TeV. *Phys. Lett.*, B743:503–525, 2015.

[3] Jiji Fan, Matthew Reece, and Joshua T. Ruderman. A stealth supersymmetry sampler. *Journal of High Energy Physics*, 2012(7), 2012.

[4] Aneesh V. Manohar and Mark B. Wise. Flavor changing neutral currents, an extended scalar sector, and the Higgs production rate at the CERN Large Hadron Collider. *Physical Review D - Particles, Fields, Gravitation and Cosmology*, 74(3), 2006.

[5] Michael Gerbush, Teng Jian Khoo, Daniel J. Phalen, Aaron Pierce, and David Tucker-Smith. Color-octet scalars at the CERN LHC. *Physical Review D - Particles, Fields, Gravitation and Cosmology*, 77(9), 2008.

[6] Jonathan M. Arnold and Bartosz Fornal. Color octet scalars and high p T four-jet events at the LHC. *Physical Review D - Particles, Fields, Gravitation and Cosmology*, 85(5), 2012.

[7] Herbi K. Dreiner, Howard E. Haber, and Stephen P. Martin. Two-component spinor techniques and Feynman rules for quantum field theory and supersymmetry. *Phys. Rept.*, 494:1–196, 2010.

[8] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro. The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations. *Journal of High Energy Physics*, 2014(7), 2014.

[9] Kfir Blum, Aielet Efrati, Claudia Frugiuele, and Yosef Nir. Exotic colored scalars at the LHC. 2016.

[10] Dilani Kahawala and Yevgeny Kats. Distinguishing spins at the LHC using bound state signals. *JHEP*, 09:099, 2011.

[11] Yevgeny Kats and Matthew J. Strassler. Probing Colored Particles with Photons, Leptons, and Jets. *JHEP*, 11:097, 2012. [Erratum: JHEP07,009(2016)].

[12] Adam Alloul, Neil D. Christensen, Céline Degrande, Claude Duhr, and Benjamin Fuks. FeynRules 2.0 - A complete toolbox for tree-level phenomenology. *Computer Physics Communications*, 185(8):2250–2300, 2014.

[13] Torbjörn Sjöstrand, Stephen Mrenna, and Peter Skands. A brief introduction to PYTHIA 8.1. *Computer Physics Communications*, 178(11):852–867, 2008.

[14] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. FastJet User Manual. *Eur. Phys. J.*, C72:1896, 2012.

[15] V. Khachatryan and Others. Search for supersymmetry in events with soft leptons, low jet multiplicity, and missing transverse energy in proton-proton collisions at s=8 TeV. *Physics Letters, Section B: Nuclear, Elementary Particle and High-Energy Physics*, 759:9–35, 2016.

[16] J. De Favereau, C. Delaere, P. Demin, A. Giammanco, V. Lemaître, A. Mertens, and M. Selvaggi. DELPHES 3: A modular framework for fast simulation of a generic collider experiment. *Journal of High Energy Physics*, 2014(2), 2014.

[17] Morad Aaboud et al. Search for new phenomena in a lepton plus high jet multiplicity final state with the ATLAS experiment using $\sqrt{s} = 13$ TeV proton-proton collision data. *JHEP*, 09:088, 2017.

[18] Georges Aad et al. A search for $t\bar{t}$ resonances using lepton-plus-jets events in proton-proton collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector. *JHEP*, 08:148, 2015.

[19] Serguei Chatrchyan et al. Search for supersymmetry in pp collisions at $\sqrt{s}$=8 TeV in events with a single lepton, large jet multiplicity, and multiple b jets. *Phys. Lett.*, B733:328–353, 2014.

[20] Morad Aaboud et al. Search for top-squark pair production in final states with one lepton, jets, and missing transverse momentum using 36 fb^{-1} of $\sqrt{s} = 13$ TeV pp collision data with the ATLAS detector. 2017.

[21] Albert M Sirunyan et al. Search for R-parity violating supersymmetry in pp collisions at $\sqrt{s} = 13$ TeV using b jets in a final state with a single lepton, many jets, and high sum of large-radius jet masses. *Submitted to: Phys. Lett. B*, 2017.

[22] Morad Aaboud et al. A search for pair-produced resonances in four-jet final states at $\sqrt{s} =$ 13 TeV with the ATLAS detector. *Eur. Phys. J.*, C78(3):250, 2018.

[23] C. Amsler et al. *Physics Letters*, B667:1, 2008 and 2009 partial update for the 2010 edition.

Chapter 8

Conclusion and Outlook

You've made it to the end, congratulations! Feel free to contact me with any questions or suggestions on this work. Most of the work contained in this thesis has been to investigate the phenomenology of various theories at beam dump and collider experiments. In doing so, we have developed simple and powerful searches for constraining new physics, demonstrated their power by leveraging decades old data to obtain new bounds, and shown how future experiments can further improve upon these bounds. We have also contributed towards the case for building the milliQan experiment by performing a full Geant4 simulation of its response to millicharged particles, thus providing a very detailed study of its sensitivity.

Although most of the focus in this thesis has been on searches for new physics, it is important to remember that the way forward towards actually finding new physics relies equally on improving our understanding of SM physics, and the methods we use in our calculations. Our simulations of HNLs, mCPs and NTP at beam dump experiments relied heavily on the following inputs: meson production spectrums as a function of POT energies, nuclei form factors, efficient large N phase space integrals, background determinations for our channels of interest, and efficiency responses of the detectors as a function of the particle type and energy. Many of this information was difficult to obtain, or simply not known. If

we are to broaden the reach of this upcoming generation of neutrino experiments to theories outside of their flagship measurements, it will be important for the phenomenology and experimental communities to join forces.

Moving forward, I believe it will be increasingly important to take an 'open source' approach so that phenomenologists can have greater insight into the capabilities of neutrino experiments. I put forward and re-emphasize some useful ideas, some of which have of course already been implemented:

- Write review papers cataloging the various meson distributions and form factors relevant at a variety of beam dump experiments.
- Publish data and analyses for various final state topologies, over a broad range of energies. These will act as backgrounds to various new searches.
- Publish particle detection efficiencies, thresholds and detector capabilities at various experiments.
- Write open source C++ libraries that are integrated with Feynrules and efficiently perform large N-body phase space integrals with cuts for fixed target processes.
- Release fast-running software to simulate the whole beam dump experiment pipeline, from the meson production and decays, to the detection and analysis. This software should also be integrated with Feynrules, should the user wish to input new physics models.
- Organize international workshops between theorists and experimentalists so that the two communities can discuss their respective problems and challenges.
- List experimental challenges that theorists should work on. For instance, combining near and far detector information to reduce systematic uncertainties on the initial neutrino fluxes and flavor contamination.

The difficulty in working on any one of these items is that they might only be relevant for a very small subset of people. However, the high energy community (as related to the LHC, ATLAS and CMS) has demonstrated that once all of these tools exist, have clear templates and examples, and are maintained by a

dedicated group, the whole community will rally around them. This has been the case with Les Houches and the MSTW collaborations for their implementations of parton distribution functions, and with FeynRules, MadGraph, Pythia, Delphes for their implementation of the whole LHC pipeline. The dialogue between phenomenologists and experimentalists in high energy has also brought much benefit to the experimental collaborations. Countless new techniques that were developed for sensitivity analyses from the theoretical side are now in use by the collaborations [96, 97, 78, 98]. On the neutrino end, the GENIE software [99] and the upcoming MadDump software (a MadGraph-like software for beam dump experiments) are great examples that address many of the suggestions made above. In light of the uncertainty towards where new physics is hiding, it makes sense to leverage the existing and upcoming wave of neutrino experiments and take full advantage of the large luminosities and state of the art technology at their disposal. In addition to focusing on new physics, encouragement and funding should be given for people to work on developing the methods listed above. Any progress in this area really helps drive the field forward, saves time, and makes the work of phenomenologists better aligned with the searches carried out by the experiments.

In conclusion, it will be very exciting to see where the future of particle physics is headed. In the high energy frontier, the LHC running at 14 TeV will deliver about 75 times more data than has been collected thus far. At the intensity frontier, there is a wave of very capable high performing near and far LAr-TPC neutrino detectors currently being built that will determine the structure of neutrinos and provide us with many new interesting measurements. And lastly, new astrophysical data from gravitational waves [100, 101], CMB [102] and 21 cm physics [103, 104, 105] are already revealing interesting findings about the universe. At any rate, it is important to leverage all of our sources of data, and be as best equipped as possible to efficiently analyze it under a variety of new physics hypotheses.

Bibliography

[1] On a New Kind of Rays. *Nature*, 53:274–276, January 1896.

[2] Georges 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.*, B716:1–29, 2012.

[3] Serguei Chatrchyan et al. Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC. *Phys. Lett.*, B716:30–61, 2012.

[4] Carsten Burgard. Standard model of physics. http://www.texample.net/tikz/examples/model-physics/. Accessed: 2018-05-24.

[5] R. Acciarri et al. Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE). 2015.

[6] R. Acciarri et al. Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE). 2016.

[7] Sergey Alekhin et al. A facility to Search for Hidden Particles at the CERN SPS: the SHiP physics case. *Rept. Prog. Phys.*, 79(12):124201, 2016.

[8] M. Anelli et al. A facility to Search for Hidden Particles (SHiP) at the CERN SPS. 2015.

[9] M. Antonello et al. A Proposal for a Three Detector Short-Baseline Neutrino Oscillation Program in the Fermilab Booster Neutrino Beam. 2015.

[10] C. P. Burgess and G. D. Moore. *The standard model: A primer*. Cambridge University Press, 2006.

[11] P. F. de Salas, D. V. Forero, C. A. Ternes, M. Tortola, and J. W. F. Valle. Status of neutrino oscillations 2018: first hint for normal mass ordering and improved CP sensitivity. 2017.

[12] Guido Fantini, Andrea Gallo Rosso, Francesco Vissani, and Vanessa Zema. Introduction to the Formalism of Neutrino Oscillations. *Adv. Ser. Direct. High Energy Phys.*, 28:37–119, 2018.

[13] A. A. Aguilar-Arevalo et al. Unexplained Excess of Electron-Like Events From a 1-GeV Neutrino Beam. *Phys. Rev. Lett.*, 102:101802, 2009.

[14] A. A. Aguilar-Arevalo et al. A Search for electron neutrino appearance at the $\Delta m^2 \sim 1\text{eV}^2$ scale. *Phys. Rev. Lett.*, 98:231801, 2007.

[15] C. Athanassopoulos et al. Evidence for anti-muon-neutrino —¿ anti-electron-neutrino oscillations from the LSND experiment at LAMPF. *Phys. Rev. Lett.*, 77:3082–3085, 1996.

[16] Joachim Kopp, Pedro A. N. Machado, Michele Maltoni, and Thomas Schwetz. Sterile Neutrino Oscillations: The Global Picture. *JHEP*, 05:050, 2013.

[17] J. M. Conrad and M. H. Shaevitz. Sterile Neutrinos: An Introduction to Experiments. *Adv. Ser. Direct. High Energy Phys.*, 28:391–442, 2018.

[18] S. N. Gninenko. The MiniBooNE anomaly and heavy neutrino decay. *Phys. Rev. Lett.*, 103:241802, 2009.

[19] Sergei N. Gninenko. A resolution of puzzles from the LSND, KARMEN, and MiniBooNE experiments. *Phys. Rev.*, D83:015015, 2011.

[20] A. A. Aguilar-Arevalo et al. Observation of a Significant Excess of Electron-Like Events in the MiniBooNE Short-Baseline Neutrino Experiment. 2018.

[21] Richard J. Hill. On the single photon background to ν_e appearance at MiniBooNE. *Phys. Rev.*, D84:017501, 2011.

[22] A. A. Aguilar-Arevalo et al. The Neutrino Flux prediction at MiniBooNE. *Phys. Rev.*, D79:072002, 2009.

[23] Emidio Gabrielli, Luca Marzola, Martti Raidal, and Hardi Veermäe. Dark matter and spin-1 milli-charged particles. *JHEP*, 08:150, 2015.

[24] Robert Foot, H. Lew, and R. R. Volkas. Electric charge quantization. *J. Phys.*, G19:361–372, 1993. [Erratum: J. Phys.G19,1067(1993)].

[25] Eder Izaguirre and Itay Yavin. New window to millicharged particles at the LHC. *Phys. Rev.*, D92(3):035014, 2015.

[26] Serguei Chatrchyan et al. Search for fractionally charged particles in pp collisions at $\sqrt{s} = 7$ TeV. *Phys. Rev.*, D87(9):092008, 2013.

[27] Andrew Haas, Christopher S. Hill, Eder Izaguirre, and Itay Yavin. Looking for milli-charged particles with a new experiment at the LHC. *Phys. Lett.*, B746:117–120, 2015.

[28] Austin Ball et al. A Letter of Intent to Install a milli-charged Particle Detector at LHC LEPP5. 2016.

[29] Gabriel Magill. A Milli-Charged Particle Detector at LHC P5. DESY-PROC, pages 39–42, Hamburg, Jun 2017. 12th Patras Workshop on Axions, WIMPs and WISPs, Jeju Island (South Korea), 20 Jun 2016 - 24 Jun 2016, Verlag Deutsches Elektronen-Synchrotron.

[30] Gabriel Magill, Ryan Plestid, Maxim Pospelov, and Yu-Dai Tsai. Dipole portal to heavy neutral leptons. 2018.

[31] Manuel Masip, Pere Masjuan, and Davide Meloni. Heavy neutrino decays at MiniBooNE. *JHEP*, 01:106, 2013.

[32] G. G. Raffelt. *Stars as laboratories for fundamental physics*. University of Chicago Press, 1996.

[33] Raj Gandhi and Adam Burrows. Massive dirac neutrinos and the sn1987a signal. *Physics Letters B*, 246(1):149 – 155, 1990.

[34] K. Hirata et al. Observation of a neutrino burst from the supernova sn1987a. *Phys. Rev. Lett.*, 58:1490–1493, Apr 1987.

[35] R. M. Bionta et al. Observation of a Neutrino Burst in Coincidence with Supernova SN 1987a in the Large Magellanic Cloud. *Phys. Rev. Lett.*, 58:1494, 1987.

[36] Jae Hyeok Chang, Rouven Essig, and Samuel D. McDermott. Revisiting Supernova 1987A Constraints on Dark Photons. *JHEP*, 01:107, 2017.

[37] H. K. Dreiner, C. Hanhart, U. Langenfeld, and Daniel R. Phillips. Supernovae and light neutralinos: SN1987A bounds on supersymmetry revisited. *Phys. Rev.*, D68:055004, 2003.

[38] Daniel Schmidt, Thomas Schwetz, and He Zhang. Status of the Zee–Babu model for neutrino mass and possible tests at a like-sign linear collider. *Nucl. Phys.*, B885:524–541, 2014.

[39] Juan Herrero-Garcia, Miguel Nebot, Nuria Rius, and Arcadi Santamaria. The Zee–Babu model revisited in the light of new data. *Nucl. Phys.*, B885:542–570, 2014.

[40] P. S. Bhupal Dev, Clara Miralles Vila, and Werner Rodejohann. Naturalness in testable type II seesaw scenarios. *Nucl. Phys.*, B921:436–453, 2017.

[41] Miguel Nebot, Josep F. Oliver, David Palao, and Arcadi Santamaria. Prospects for the Zee-Babu Model at the CERN LHC and low energy experiments. *Phys. Rev.*, D77:093013, 2008.

[42] Junjie Cao, Peihua Wan, Lei Wu, and Jin Min Yang. Lepton-Specific Two-Higgs Doublet Model: Experimental Constraints and Implication on Higgs Phenomenology. *Phys. Rev.*, D80:071701, 2009.

[43] Wolfgang Altmannshofer, Stefania Gori, Maxim Pospelov, and Itay Yavin. Neutrino Trident Production: A Powerful Probe of New Physics with Neutrino Beams. *Phys. Rev. Lett.*, 113:091801, 2014.

[44] Gabriel Magill and Ryan Plestid. Probing new charged scalars with neutrino trident production. *Phys. Rev.*, D97(5):055003, 2018.

[45] Gabriel Magill and Ryan Plestid. Neutrino Trident Production at the Intensity Frontier. *Phys. Rev.*, D95(7):073004, 2017.

[46] S. R. Mishra et al. Neutrino tridents and W Z interference. *Phys. Rev. Lett.*, 66:3117–3120, 1991.

[47] T. Adams et al. Evidence for diffractive charm production in muon-neutrino Fe and anti-muon-neutrino Fe scattering at the Tevatron. *Phys. Rev.*, D61:092001, 2000.

[48] D. Geiregat et al. First observation of neutrino trident production. *Phys. Lett.*, B245:271–275, 1990.

[49] C. Athanassopoulos et al. The Liquid scintillator neutrino detector and LAMPF neutrino source. *Nucl. Instrum. Meth.*, A388:149–172, 1997.

[50] A. A. Aguilar-Arevalo et al. The MiniBooNE Detector. *Nucl. Instrum. Meth.*, A599:28–46, 2009.

[51] Mitchell Soderberg. MicroBooNE: A New Liquid Argon Time Projection Chamber Experiment. *AIP Conf. Proc.*, 1189:83–87, 2009.

[52] Nicola McConkey. The LAr1-ND Experiment. *J. Phys. Conf. Ser.*, 650(1):012007, 2015.

[53] A. M. Sirunyan et al. Particle-flow reconstruction and global event description with the CMS detector. *JINST*, 12(10):P10003, 2017.

[54] G. Aad et al. Expected Performance of the ATLAS Experiment - Detector, Trigger and Physics. 2009.

[55] B. Adeva et al. The construction of the l3 experiment. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 289(1):35 – 102, 1990.

[56] K Ahmet et al. The OPAL detector at LEP. *Phys. Lett. B*, 251(CERN-PPE-90-114):211–222. 93 p, Aug 1990. Contribution to the Festschrift : M Morpurgo.

[57] The delphi detector at lep. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 303(2):233 – 276, 1991.

[58] C. Patrignani et al. Review of Particle Physics. *Chin. Phys.*, C40(10):100001, 2016.

[59] L. Wasserman and L.A. Wasserman. *All of Statistics: A Concise Course in Statistical Inference*. Springer Texts in Statistics. Springer, 2004.

[60] Kyle Cranmer. Practical Statistics for the LHC. In *Proceedings, 2011 European School of High-Energy Physics (ESHEP 2011): Cheile Gradistei, Romania, September 7-20, 2011*, pages 267–308, 2015. [,247(2015)].

[61] S. S. Wilks. The large-sample distribution of the likelihood ratio for testing composite hypotheses. *The Annals of Mathematical Statistics*, 9(1):60–62, 1938.

[62] U. D. Jentschura and V. G. Serbo. Nuclear form factor, validity of the equivalent photon approximation and Coulomb corrections to muon pair production in photon-nucleus and nucleus-nucleus collisions. *Eur. Phys. J.*, C64:309–317, 2009.

[63] C. F. Perdrisat, V. Punjabi, and M. Vanderhaeghen. Nucleon Electromagnetic Form Factors. *Prog. Part. Nucl. Phys.*, 59:694–764, 2007.

[64] Douglas H. Beck and Barry R. Holstein. Nucleon structure and parity violating electron scattering. *Int. J. Mod. Phys.*, E10:1–41, 2001.

[65] A. D. Martin, W. J. Stirling, R. S. Thorne, and G. Watt. Parton distributions for the LHC. *Eur. Phys. J.*, C63:189–285, 2009.

[66] C. Grupen and B. Shwartz. *Particle Detectors*. Cambridge Monographs on Particle Physics, Nuclear Physics and Cosmology. Cambridge University Press, 2008.

[67] Janet M. Conrad, William C. Louis, and Michael H. Shaevitz. The LSND and MiniBooNE Oscillation Searches at High Δm^2. *Ann. Rev. Nucl. Part. Sci.*, 63:45–67, 2013.

[68] O. Lalakulich, U. Mosel, and K. Gallmeister. Neutrino energy reconstruction in quasielastic-like scattering in the miniboone and t2k experiments. *Phys. Rev. C*, 86:054606, Nov 2012.

[69] L. Alvarez-Ruso et al. NuSTEC 1 1 Neutrino Scattering Theory Experiment Collaboration http://nustec.fnal.gov . White Paper: Status and challenges of neutrino–nucleus scattering. *Prog. Part. Nucl. Phys.*, 100:1–68, 2018.

[70] R.L. Burman and E.S. Smith. Parameterization of pion production and reaction cross sections at lampf energies. *LAMPF Report*, LA-11502-MS, 5 1989.

[71] Sanford J. R. and C.L. Wang. Empirical formulas for particle production in p-be collision between 10 and 35 bev/c part i. 3 1967.

[72] C. Mariani, G. Cheng, J. M. Conrad, and M. H. Shaevitz. Improved Parameterization of K^+ Production in p-Be Collisions at Low Energy Using Feynman Scaling. *Phys. Rev.*, D84:114021, 2011.

[73] M. Bonesini, A. Marchionni, F. Pietropaolo, and T. Tabarelli de Fatis. On Particle production for high-energy neutrino beams. *Eur. Phys. J.*, C20:13–27, 2001.

[74] Charles Gale, Sangyong Jeon, and Joseph I. Kapusta. J / psi production and absorption in high-energy proton nucleus collisions. *Phys. Lett.*, B459:455–460, 1999.

[75] Georges Aad et al. Search for supersymmetry at $\sqrt{s} = 13$ TeV in final states with jets and two same-sign leptons or three leptons with the ATLAS detector. *Eur. Phys. J.*, C76(5):259, 2016.

[76] Vardan Khachatryan et al. Search for supersymmetry in events with a photon, a lepton, and missing transverse momentum in pp collisions at $\sqrt{s} = 8$ TeV. *Phys. Lett.*, B757:6–31, 2016.

[77] Torbjörn Sjöstrand, Stefan Ask, Jesper R. Christiansen, Richard Corke, Nishita Desai, Philip Ilten, Stephen Mrenna, Stefan Prestel, Christine O. Rasmussen, and Peter Z. Skands. An Introduction to PYTHIA 8.2. *Comput. Phys. Commun.*, 191:159–177, 2015.

[78] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. The Anti-k(t) jet clustering algorithm. *JHEP*, 04:063, 2008.

[79] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro. The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations. *JHEP*, 07:079, 2014.

[80] Adam Alloul, Neil D. Christensen, Celine Degrande, Claude Duhr, and Benjamin Fuks. FeynRules 2.0 - A complete toolbox for tree-level phenomenology. *Comput. Phys. Commun.*, 185:2250–2300, 2014.

[81] Neil D. Christensen and Claude Duhr. FeynRules - Feynman rules made easy. *Comput. Phys. Commun.*, 180:1614–1641, 2009.

[82] G. Corcella, I. G. Knowles, G. Marchesini, S. Moretti, K. Odagiri, P. Richardson, M. H. Seymour, and B. R. Webber. HERWIG 6: An Event generator for hadron emission reactions with interfering gluons (including supersymmetric processes). *JHEP*, 01:010, 2001.

[83] T. Gleisberg, Stefan. Hoeche, F. Krauss, M. Schonherr, S. Schumann, F. Siegert, and J. Winter. Event generation with SHERPA 1.1. *JHEP*, 02:007, 2009.

[84] S. Agostinelli et al. GEANT4: A Simulation toolkit. *Nucl. Instrum. Meth.*, A506:250–303, 2003.

[85] I. Antcheva et al. ROOT: A C++ framework for petabyte data storage, statistical analysis and visualization. *Comput. Phys. Commun.*, 180:2499–2512, 2009.

[86] A. A. Prinz et al. Search for millicharged particles at SLAC. *Phys. Rev. Lett.*, 81:1175–1178, 1998.

[87] L. B. Auerbach et al. Measurement of electron - neutrino - electron elastic scattering. *Phys. Rev.*, D63:112001, 2001.

[88] Bob Holdom. Two U(1)'s and Epsilon Charge Shifts. *Phys. Lett.*, 166B:196–198, 1986.

[89] A. D. Dolgov, S. L. Dubovsky, G. I. Rubtsov, and I. I. Tkachev. Constraints on millicharged particles from Planck data. *Phys. Rev.*, D88(11):117701, 2013.

[90] Sacha Davidson, Steen Hannestad, and Georg Raffelt. Updated bounds on milli-charged particles. *Journal of High Energy Physics*, 2000(05):003, 2000.

[91] A. Badertscher, P. Crivelli, W. Fetscher, U. Gendotti, S. Gninenko, V. Postoev, A. Rubbia, V. Samoylenko, and D. Sillou. An Improved Limit on Invisible Decays of Positronium. *Phys. Rev.*, D75:032004, 2007.

[92] Hamamatsu. photomultiplier tube r329-02. http://www.hamamatsu.com/us/en/product/category/3100/3001/R329-02/index.html. Accessed: 2016-02-12.

[93] Saint-gobain. plastic scintillators. http://www.crystals.saint-gobain.com/Plastic_Scintillators.aspx. Accessed: 2016-02-12.

[94] Vardan Khachatryan et al. Search for pair-produced resonances decaying to jet pairs in proton–proton collisions at $\sqrt{s} = 8$ TeV. *Phys. Lett.*, B747:98–119, 2015.

[95] Morad Aaboud et al. A search for pair-produced resonances in four-jet final states at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Eur. Phys. J.*, C78(3):250, 2018.

[96] Jesse Thaler and Ken Van Tilburg. Identifying Boosted Objects with N-subjettiness. *JHEP*, 03:015, 2011.

[97] Eder Izaguirre, Brian Shuve, and Itay Yavin. Improving Identification of Dijet Resonances at Hadron Colliders. *Phys. Rev. Lett.*, 114(4):041802, 2015.

[98] Tilman Plehn, Gavin P. Salam, and Michael Spannowsky. Fat Jets for a Light Higgs. *Phys. Rev. Lett.*, 104:111801, 2010.

[99] C. Andreopoulos et al. The GENIE Neutrino Monte Carlo Generator. *Nucl. Instrum. Meth.*, A614:87–104, 2010.

[100] B.P. Abbott et al. GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral. *Phys. Rev. Lett.*, 119(16):161101, 2017.

[101] B.P. Abbott et al. Observation of Gravitational Waves from a Binary Black Hole Merger. *Phys. Rev. Lett.*, 116(6):061102, 2016.

[102] P. A. R. Ade et al. Planck 2015 results. XIII. Cosmological parameters. *Astron. Astrophys.*, 594:A13, 2016.

[103] Judd D. Bowman, Alan E. E. Rogers, Raul A. Monsalve, Thomas J. Mozdzen, and Nivedita Mahesh. An absorption profile centred at 78 megahertz in the sky-averaged spectrum. *Nature*, 555:67 EP –, Feb 2018.

[104] Rennan Barkana. Possible interaction between baryons and dark-matter particles revealed by the first stars. *Nature*, 555:71 EP –, Feb 2018.

[105] Asher Berlin, Dan Hooper, Gordan Krnjaic, and Samuel D. McDermott. Severely Constraining Dark Matter Interpretations of the 21-cm Anomaly. 2018.