

Dipole portal to heavy neutral leptons

Gabriel Magill,^{1,2,†} Ryan Plestid,^{1,2,‡} Maxim Pospelov,^{1,3,4,§} and Yu-Dai Tsai^{5,6,*}

¹*Perimeter Institute for Theoretical Physics, 31 Caroline Street North,
Waterloo, Ontario N2L 2Y5, Canada*

²*Department of Physics & Astronomy, McMaster University, 1280 Main Street West,
Hamilton, Ontario L8S 4M1, Canada*

³*Department of Physics and Astronomy, University of Victoria, Victoria,
British Columbia V8P 5C2, Canada*

⁴*Theoretical Physics Department, CERN, 1211 Geneva, Switzerland*

⁵*Fermilab, Fermi National Accelerator Laboratory, Batavia, Illinois 60510, USA*

⁶*University of Chicago, Kavli Institute for Cosmological Physics, Chicago, Illinois 60637, USA*



(Received 27 March 2018; published 11 December 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.

DOI: [10.1103/PhysRevD.98.115015](https://doi.org/10.1103/PhysRevD.98.115015)

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 (HNL), 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,

$$\mathcal{L} \supset \bar{N}(i\not{\partial} - m_N)N + (d\bar{\nu}_L\sigma_{\mu\nu}F^{\mu\nu}N + \text{H.c.}) \quad (1)$$

Here $F^{\mu\nu}$ is the electromagnetic field strength tensor, and ν_L is a SM neutrino field. This is an effective Lagrangian that

*Corresponding author.
ytsai@fnal.gov, yt444@cornell.edu
†gmagill@perimeterinstitute.ca
‡plestird@mcmaster.ca
§mpospelov@perimeterinstitute.ca

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by SCOAP³.

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, 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 $N \rightarrow \nu\gamma$ ([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 quasielastic events by MiniBoone’s Cherenkov detector.)

The most popular proposal is the existence of a light ($m \sim \text{eV}$) 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 (d.o.f.) in the early Universe, but also limits the total neutrino mass $\sum m_\nu \leq 0.17 \text{ eV}$ 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

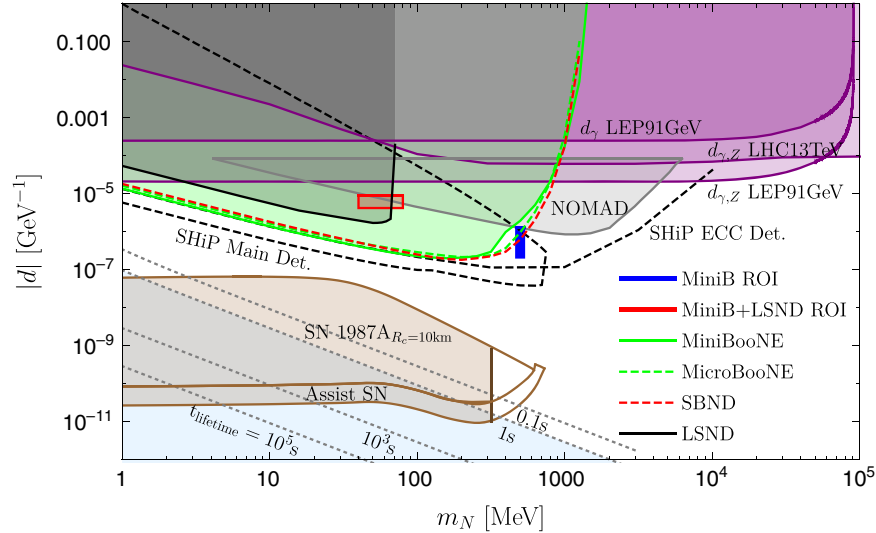


FIG. 1. Overview of projected sensitivities (95% CL) and constraints obtained from SHiP, LHC, LEP, Supernova 1987A and experiments at the Short-Baseline Neutrino facility at Fermilab. We also show previously calculated favored regions of interest (ROI) in parameter space for MiniBooNE and LSND, and constraints 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.

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 midrange 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 s–1 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 $SU(2) \times U(1)$ 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 Sec. II, we provide more details on the model including the possible SM gauge invariant completions and the connections to neutrino masses. In Secs. III–V, we consider the intensity, energy and astrophysical frontiers respectively. Finally, we conclude in Sec. VI 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




FIG. 2. Dipole portal processes for N : (a) Production of N from off-shell neutrinos arising from weak meson decays (e.g., from π , $K \rightarrow \mu^+ \nu^*$); (b) Production of N from off-shell photons arising from Dalitz-like meson decays (e.g., from π^0 , $\eta \rightarrow \gamma^* \gamma$); (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.

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.

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 $-\frac{1}{R_{\text{nuc}}^2} \leq t \leq -\frac{m_N^2}{4E_\nu^2}$, we find

$$\sigma_{\nu \rightarrow N}^E = 4\alpha Z^2 |d|^2 \times \log\left(\frac{4E_\nu^2}{m_N^4 R_{\text{nuc}}^2}\right), \quad (2)$$

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 $\nu \rightarrow N$ upscattering due to the magnetic moment of the nucleus. Cutting off the coherent scattering at $|t| = \frac{1}{R_{\text{nuc}}^2}$, one has

$$\sigma_{\nu \rightarrow N}^M = \frac{(I+1)}{3\pi I} \mu_A^2 |d|^2 \times \frac{1}{R_{\text{nuc}}^2} < \sigma_E/Z^2, \quad (3)$$

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 $\nu \rightarrow N$ 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 $R_{\text{nuc}}^{-2} \sim 0.3$ GeV², we can expect a production cross section per nucleus of roughly

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

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 $\text{Br}_{M \rightarrow N} \propto d^2 m_M^2$, where m_M is the mass of the decaying meson.

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

$$\Gamma_{N \rightarrow \nu \gamma} = \frac{|d|^2 m_N^3}{4\pi}, \quad (5)$$

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

$$t_{\text{dec}} = \tau_\gamma = 1.3 \times 10^{-6} \text{ s} \left(\frac{50 \text{ MeV}}{m_N}\right)^4 \left(\frac{10^{-6} \text{ GeV}^{-1}}{|d|}\right)^2$$

$$L_{\text{dec}} = c\tau_\beta \gamma \approx 400 \text{ m} \left(\frac{50 \text{ MeV}}{m_N}\right)^4 \left(\frac{10^{-6} \text{ GeV}^{-1}}{|d|}\right)^2. \quad (6)$$

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

$$N_D = \begin{pmatrix} N \\ N^{c\dagger} \end{pmatrix}, \quad (7)$$

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, \mathbf{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, \mathbf{m}_ν will scale as $\mathbf{m}_\nu \sim d^2 \Lambda^2 \mathbf{m}_N / 16\pi^2$, 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 $\mathbf{m}_N \ll m_N$. Quasi-Dirac N would typically lead to *larger* values of d than otherwise would be suggested by a simple application of the seesaw 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

$$d = \frac{3m_{\nu N} e G_F}{32\pi^2 \sqrt{2}} = 1.2 \times 10^{-9} \text{ GeV}^{-1} \left(\frac{m_{\nu N}}{50 \text{ MeV}} \right). \quad (8)$$

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 $m_N = 50 \text{ MeV}$, observed neutrino masses would imply $m_{\nu N} \sim \text{keV}$ and consequently $d \sim 10^{-13} \text{ GeV}^{-1}$. We do not impose such a stringent constraint and consider d to be an independent




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.

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 $SU(2) \times U(1)$ interpretation of d would require a Higgs insertion, so that the dipole interaction is really a dimension 6 operator. Therefore, in the limit of large Λ the maximum expected d is

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

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 $d_{\text{max,pert}} \sim \text{GeV}/\Lambda^2$. To consider neutrino dipole couplings which respect the full gauge symmetries of the standard model, we write down the Lagrangian

$$\mathcal{L} \supset \bar{L}(d_W \mathcal{W}_{\mu\nu}^a \tau^a + d_B \mathcal{B}_{\mu\nu}) \tilde{H} \sigma_{\mu\nu} N_D + \text{H.c.} \quad (10)$$

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

$$\mathcal{L} \supset d_W (\bar{\ell}_L W_{\mu\nu}^- \sigma^{\mu\nu} N_D) + \bar{\nu}_L [d_\gamma F_{\mu\nu} - d_Z Z_{\mu\nu}] \sigma^{\mu\nu} N_D + \text{H.c.} \quad (11)$$

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

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

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_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_B^e and d_B^μ . 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

$d_{\gamma,B,W}^\mu$ coupling for simplicity. One can also turn on $d_{\gamma,B,W}^e$ and $d_{\gamma,B,W}^\tau$ 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 Sec. III B. Reference [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 $SU(3) \times SU(2) \times U(1)$). 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 $2 \rightarrow N$, 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 reabsorbed 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 $\Delta m^2 \sim 1 \text{ eV}^2$. It will consist of 3 LAr-TPC detectors called SBND, MicroBooNE and ICARUS, which will be spread out over a 600 m 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 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

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




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.

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) $\text{BR}(\text{Meson} \rightarrow X + \gamma) = \mathcal{O}(1)$ (e.g., π^0 , η) or $\text{BR}(\text{Meson} \rightarrow X + \nu) = \mathcal{O}(1)$ (e.g., π^\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 50 m meson decay chamber, whereas SHiP has a hadron stopper. This divides our discussion into prompt ($\tau^{\text{rest}} \lesssim 10^{-12}$ s) and long-lived ($\tau^{\text{rest}} \gtrsim 10^{-12}$ 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

$$\pi^0, \eta \rightarrow \gamma(\gamma^* \rightarrow \nu_a N) \quad (14)$$

We immediately see that these radiative Dalitz-like decays will be useful for improving the sensitivity to d_e and d_τ

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.

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

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 π^\pm and K^\pm . They can produce HNLs via an off-shell neutrino decay

$$\pi^\pm, K^\pm \rightarrow \mu^\pm(\nu_\mu^{(-)*} \rightarrow \gamma N). \quad (15)$$

When considering decays to electron flavor, such as $K, \pi \rightarrow e\nu_e$, 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 $K, \pi \rightarrow eN\gamma$ 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. 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 misreconstructed π^0 or radiative resonance decays such as $\Delta \rightarrow N\gamma$. 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 electronlike 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 Figs. 1 and 10 of [33], we note that the incident neutrino flux favors energy values between 30–50 MeV, whereas the collected electronlike 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 data set 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 $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ 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 modeling 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 modeled 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 $\bar{\nu}_\mu$, 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 $L_{\text{dec}} \gg L_{\text{fid}}$, 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 $\nu_\mu \rightarrow \nu_e$ quasielastic 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 $\Delta \rightarrow N\gamma$ 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. First, we use the data and backgrounds as given in [17]. Second, we repeat the analysis after including the additional sources of backgrounds identified in [18]. And last, 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




FIG. 5. 95% CL limits for HNL particles using MiniBooNE and LSND ν_e CC 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 $\Delta \rightarrow \gamma N$ 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).

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 $N \rightarrow \nu_{H'}\gamma$ 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].

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.

3. SHiP and NOMAD

At CERN, we will consider the NOMAD experiment, which ran from 1995 to 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




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).

single photon production. Using this data corresponding to 1.45×10^{18} POT, Monte Carlo simulations of HNL signals (with no mass mixing) were performed [44,45] to simulate the Primakoff process $\nu_\mu Z \rightarrow NZ$. 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 results¹ 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 100 m. 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.

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

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




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.

on d_f^f , for a given flavor f . Recall that flavor indices in Eq. (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




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.

assume 1000 background events in both detectors in order to provide a conservative estimate.

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 and 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 $d_{B,\mathcal{W}} \equiv \sqrt{2} \overline{d_{B,\mathcal{W}}} / (v \cos \theta_w)$ to obtain

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

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 $e^+e^- \rightarrow (N \rightarrow \gamma\nu)\bar{\nu} + \text{H.c.}$ 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 $e^-e^+ \rightarrow N\bar{\nu}$ integrated over all angles is

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 $SU(2) \times U(1)$ symmetry of the standard model.

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	$d_{\gamma,Z}$	$\overline{d_{\mathcal{W}}} = 0$	$\overline{d_B}$
	$d_{\gamma,W}^a$	$d_\gamma = a \times \overline{d_{\mathcal{W}}}$	$\overline{d_{\mathcal{W}}}$

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

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 Sec. IV B and Appendix C and do not make approximations on the masses of electrons.

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 anti-quarks interacting via an s-channel photon or Z boson: $q_i \bar{q}_i \rightarrow (N \rightarrow \gamma\nu)\bar{\nu} + \text{H.c.}$ This gives the same signature as LEP, up to subtleties that will be discussed in Sec. 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: $u_i \bar{d}_j \rightarrow (N \rightarrow \gamma\nu)\ell^+$. 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 $\gamma + E_{\text{miss}}$ final state [50–53]. We choose to focus on the results of LEP1, which ran at a center of mass (COM) energy corresponding to the Z pole and accumulated about 200 pb⁻¹ of data, and LEP161 which ran at a COM energy of 161 GeV and accumulated 25 pb⁻¹ [54]. Using partial luminosity and combining many analyses, LEP1 was able to set an upper bound of 0.1 pb on the cross section of new physics contributing to the $\gamma + E_{\text{miss}}$ 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 1 m 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 1 pb 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 modeling 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 $Z(\rightarrow\nu\nu)\gamma$ process, followed by $W(\rightarrow\ell\nu)\gamma$ 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 1.29 m, 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 $d_{\gamma W}^a$ coupling, and decays of N to a neutrino and photon via the d_γ coupling. We do this for various relative magnitudes between $d_{\gamma W}^a$ 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




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 ($d_{\gamma W}^a$) and decay (d_γ) couplings. See Table II for an explanation of the plot labels.

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 $\ell + \gamma + \cancel{E}_T$ 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

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 ^4He 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 sec, will distort this balance, and contribute to the Hubble rate during the proton-neutron freeze-out. Since most of the neutrons end up in ${}^4\text{He}$, 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 $\nu + \gamma \rightarrow N$. These processes are important at $T \sim m_N$, and can be derived from the width of N .
- (ii) $2 \rightarrow 2$ processes, such as $f^+ f^- \rightarrow N \bar{\nu}$ 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

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

where n_f and 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

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

where M_{Pl} is the Planck mass and g_* is the effective number of d.o.f. appearing from the definition of the Hubble rate, $H(T) \simeq 1.66 g_*^{1/2} T^2 M_{\text{Pl}}^{-1}$. The most important feature of Eq. (19), besides the self-explanatory dependence on M_{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

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

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

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

where we reexpressed d in terms of the lifetime formula for N . The decoupling of N means that at temperatures $T < T_{\text{decouple}}$ 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 ρ_N/ρ_{SM} will gain an enhancement factor m_N/T_{dec} , where T_{dec} is the temperature corresponding to the time of the decay of N , $H(T_{\text{dec}}) \sim \tau_N^{-1}$ (in the assumption that $T_{\text{dec}} < m_N$). Consequently, our estimate becomes

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

where $g_N = 7/8 \times 4$ as N carries four fermionic d.o.f. This estimate can be used to constrain the lifetime of HNLs as ρ_N/ρ_{SM} is constrained at $T \sim 1$ MeV through the n/p freeze-out. If $\tau_N \sim 0.1$ 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 $T(R_\nu) = 3 \text{ MeV}$) [66], then the energy stored in the HNLs can be efficiently recycled and reemitted 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 study³ of the following integrals [66]

$$\frac{dE}{dt} = \int_0^{R_\nu} P_{\text{esc}}(r_0) \times \left\langle E_N \frac{d\Gamma_{\text{prod}}}{dr} \right\rangle (r_0) dr_0 \quad (22a)$$

$$P_{\text{esc}}(r_0) = \exp \left[- \int_{r_0}^{R_{\text{far}}} \frac{1}{\lambda_{\text{MFP}}} dr \right]. \quad (22b)$$

where $d\Gamma/dr$ is the local rate of production of HNLs, E_N denotes the HNL energy, R_{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_{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 $\lambda_{\text{MFP}} \propto 1/d^2$.

For each HNL mass m_N , there will exist a minimal dipole coupling $d_{\text{prod}}(m_N)$ for which too few HNLs are produced to significantly alter the observed neutrino signal. Likewise, there will also exist a maximum dipole coupling $d_{\text{abs}}(m_N)$ 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 these two curves in the $d - m_N$ plane *i.e.*, $d_{\text{prod}}(m_N) < d_{\text{excl}}(m_N) < d_{\text{abs}}(m_N)$. Although Eqs. (22a) and (22b) are in general complicated, in the weak coupling regime ($d \lesssim d_{\text{prod}}$), and the strong coupling regime ($d \gtrsim d_{\text{abs}}$), the analysis simplifies.

In trying to obtain the lower curve $d_{\text{prod}}(m_N)$ 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

³Equation (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].

terms of the energy carried by HNLs per unit volume, per unit time, $d\mathcal{E}_N/dt$ (being referred to as emissivity throughout this paper), at a fixed radius r_0

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

where $d\mathcal{E}_\nu/dt$ 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 $d_{\text{abs}}(m_N) \gg d_{\text{prod}}(m_N)$ 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 $P_{\text{esc}}(d_{\text{abs}}) = 1/2$, since for $d \gtrsim d_{\text{abs}}$ 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_{abs} . This is because, being produced in the hot and dense interior of the supernova, 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 $r_0^{(a)} = 10 \text{ km}$ and $r_0^{(b)} = 14 \text{ km}$. The former corresponds to the conventional choice [61,62,64,66] of the hottest ($T \approx 30 \text{ MeV}$) 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 ($T \approx 20 \text{ MeV}$) 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 $18 M_\odot$ 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 antiparticles. 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

$$\nu + e^\pm \rightarrow N + e^\pm \quad (\text{upscattering}) \quad (24)$$

$$\nu + p \rightarrow N + p \quad (\text{upscattering}) \quad (25)$$

$$e^+ + e^- \rightarrow \bar{\nu} + N \quad (\text{synthesis}) \quad (26)$$

$$\gamma + \nu \rightarrow N \quad (\text{inverse decay}). \quad (27)$$

We point out that our analysis does not include thermal field theory effects, and so we omit the ‘‘plasmon decay’’ $\gamma \rightarrow \bar{\nu}N$ 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 $e^- \nu \rightarrow e^- N$, 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 $\nu + n \rightarrow N + n$ 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

$$\frac{d\mathcal{E}_N}{dt} = \int \frac{d^3 p_1}{(2\pi)^3} \frac{d^3 p_2}{(2\pi)^3} f_1 f_2 \langle E_N \sigma \rangle_{\mathfrak{F}} v_{\text{Møll}}, \quad (28)$$

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

$$v_{\text{Møll}} = \sqrt{(\mathbf{v}_1 - \mathbf{v}_2)^2 - (\mathbf{v}_1 \times \mathbf{v}_2)^2}. \quad (29)$$

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 $2 \rightarrow 2$ 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 $2 \rightarrow 2$ processes the average is defined as

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

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

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

where

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

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 $\cos \theta \rightarrow \pm 1$ 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 $\tilde{\lambda}$, defined at a reference value $d = 10^{-7} \text{ GeV}^{-1}$ via

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

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 $\tilde{\lambda}_{\text{MFP}} \rightarrow \tilde{\lambda}_{\text{MFP}}/3$. We may then define the critical dipole moment where HNLs are efficiently trapped via the condition

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

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 $d_{\text{abs}} \rightarrow d_{\text{abs}} \times \log\left(\frac{d_{\text{abs}}}{d_{\text{prod}}}\right)^2$, 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 $2 \rightarrow 2$ scattering contribute to the mean free path. In no particular order, the relevant processes are

$$N + e^{\pm} \rightarrow \nu + e^{\pm} \quad (\text{downscattering}) \quad (35)$$

$$N + p \rightarrow \nu + p \quad (\text{downscattering}) \quad (36)$$

$$\bar{\nu} + N \rightarrow e^+ + e^- \quad (\text{annihilation}) \quad (37)$$

$$N \rightarrow \gamma + \nu \quad (\text{decay}) \quad (38)$$

$$N + \text{SN} \rightarrow N + \text{SN} \quad (\text{gravitational trapping}). \quad (39)$$

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




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

temperature within the interior of the supernova, therefore for HNLs produced at $r_0 \approx 10 \text{ km}$, 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

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

Here $\alpha \in \{e^-, e^+, \nu_{e,\mu,\tau}, \bar{\nu}_{e,\mu,\tau}, \gamma, p\}$ 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 $2 \rightarrow 2$ absorption, and the associated Pauli-blocking of outgoing SM particles for both decays and $2 \rightarrow 2$ 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 behavior 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

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

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 $|\mathcal{M}_{\text{abs}}|^2$ 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

$$\langle E_{\text{kin}} \rangle_{\text{HNL}} \leq \frac{GM_c m_N}{R_c}, \quad (42)$$

where $\langle E_{\text{kin}} \rangle_{\text{HNL}}$ 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 $M_c \approx M_{\text{SN}}$ which is the mass of SN 1987A and calculate $\langle E_{\text{kin}} \rangle_{\text{HNL}}$ at two radii $R_c = 10$ 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 $R_c = 10$ and 14 km.

C. Results

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 (^4He , 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_{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 ρ_N/ρ_{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 $\rho_N/\rho_{\text{SM}} > 0.2$ 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 $g_*(T_{\text{decouple}})$. However, m_N/T_{dec} more than compensate for this dilution,




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’’.

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 $\int \lambda_{\text{MFP}}^{-1} dr < 2/3$, 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 d.o.f., 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 Sec. 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 sec.

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 sec 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].

ACKNOWLEDGMENTS

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. Y. T. 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.,

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

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

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

The hadronic current is given by

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

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

$$W^{\mu\nu} = F_1^2 \text{Tr}[(\not{p}_4 + m_H) \gamma^\mu (\not{p}_2 + m_H) \gamma^\nu]. \quad (\text{A4})$$

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 $M_H = AM_{\text{nucleon}}$ 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 $\nu \rightarrow N$ 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 parametrizes the charge density of the nucleus as

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

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

$$\begin{aligned} \frac{d\sigma}{dt} &= -\frac{2\alpha|d|^2 Z^2 F_{\text{WS}}^2}{t^2 (s - m_H^2)^2} \\ &\times \{-tm_N^2(2s + t) + m_N^4(2m_H^2 + t) \\ &+ 2t(s - m_H^2)(-m_H^2 + s + t)\}. \end{aligned} \quad (\text{A6})$$

The $1/t^2$ prefactor 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, $M_H = M_{\text{nucleon}}$ 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

$$\begin{aligned} G_{\gamma,E}^{p,n} &= F_{1,\gamma}^{p,n} - \frac{Q^2}{4M_{\text{Nucleon}}^2} F_{2,\gamma}^{p,n} \\ G_{\gamma,M}^{p,n} &= F_{1,\gamma}^{p,n} + F_{2,\gamma}^{p,n} \end{aligned} \quad (\text{A7})$$

with

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

We then obtain

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

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 $t \rightarrow 0$ 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 $d\Gamma/d\mathcal{E}$. Normalizing by the overall decay rate of the meson defines the differential branching ratio in the rest frame $d\text{BR}/d\mathcal{E}|_{\text{rest}} = (1/\Gamma) \cdot d\Gamma/d\mathcal{E}|_{\text{rest}}$. The most important contribution to HNL production is from pions, and so we quote the result of $d\Gamma/d\mathcal{E}$ in the rest frame for the process $\pi^0 \rightarrow N\nu\gamma$

$$\begin{aligned} \frac{d\Gamma}{d\mathcal{E}} &= -\frac{1}{2\pi m_\pi} \alpha^2 |d|^2 F_\pi^2 \left[\mathcal{P} (4\mathcal{E}^2 m_\pi^2 - 3\mathcal{E} m_\pi m_N^2 \right. \\ &\quad \left. - 2\mathcal{E} m_\pi^3 + m_N^4 + 3m_\pi^2 m_N^2) \right. \\ &\quad \left. - m_\pi^2 (4\mathcal{E} - m_\pi) m_N^2 \tanh^{-1} \left(\frac{\mathcal{P}}{\mathcal{E}} \right) \right]. \end{aligned} \quad (\text{A10})$$

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 $\pi \rightarrow \mu N \gamma$ 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

$$E = \gamma\mathcal{E} - \beta\gamma\mathcal{P} \cos \phi \quad (\text{A11})$$

and noting that the decay of a pseudoscalar 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

$$\mathcal{E}_\pm = \gamma E \pm \beta\gamma P. \quad (\text{A12})$$

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

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

where the factor of $2\gamma\beta\mathcal{P}$ accounts for the normalization of the box distribution discussed above. The quantities \mathcal{E}_A and \mathcal{E}_B are defined via $\mathcal{E}_A(E, \gamma) = \min(\mathcal{E}_-, \mathcal{E}_{\min})$ and $\mathcal{E}_B(E, \gamma) = \max(\mathcal{E}_+, \mathcal{E}_{\max})$ where \mathcal{E}_{\min} and \mathcal{E}_{\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

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

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 $x_F = p_{||}^{\text{COM}}/p_{||}^{\text{COM,max}}$, 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 $A\nu \rightarrow A\nu\gamma$ 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 $\gamma\nu \rightarrow \gamma\nu$ has been explicitly calculated using effective operators [94,95], and this provides a convenient way to calculate the SM contribution to $A\nu \rightarrow A\nu\gamma$ 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 Refs. [96,97] for a comprehensive review). As discussed in Refs. [87,98] the full $\sigma_{\nu A}$ cross section can be calculated from the $\sigma_{\gamma\nu}$ cross section via

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

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 quasireal 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 $s \rightarrow 0$. 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

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

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

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

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

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

this can be rewritten as

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

Solving for s_{up} gives us the number of signal events consistent with the observation and background prediction at $(1 - \alpha)\text{CL}$. 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 underfluctuations of the data compared to the background predictions. This consists in defining

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

and solving for s_{up} in

$$\alpha' \equiv \frac{\alpha}{1 - \alpha_b} = 5\%, \quad (\text{B5})$$

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 underfluctuates. In all these cases, once we have obtained s_{up} , we can solve for the new physics coupling in the equation

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

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 $\nu(p_1)A(p_2) \rightarrow N(p_3)A'(p_4)$, followed by $N(p_3) \rightarrow \gamma(q_2)\nu(q_1)$. 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 $-0.5 \text{ GeV}^2 < t < 0$. For inelastic scattering, we limit ourselves to $-2 \text{ GeV}^2 < t < -(0.217 \text{ GeV})^2/A_3^2$, and the $t < -2 \text{ GeV}^2$ 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

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

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

$$E_N = E_\nu + \frac{t}{2M_H}. \quad (\text{C2})$$




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

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 $N \rightarrow \gamma \nu$ 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

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

We will typically choose $\theta_{\text{lab}} < \pi/4$, 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.

[1] T. Asaka and M. Shaposhnikov, *Phys. Lett. B* **620**, 17 (2005).
 [2] D. Gorbunov and M. Shaposhnikov, *J. High Energy Phys.* **10** (2007) 015; **11** (2013) 101(E).
 [3] S. Alekhin *et al.*, *Rep. Prog. Phys.* **79**, 124201 (2016).
 [4] S. N. Gninenko, *Phys. Rev. Lett.* **103**, 241802 (2009).
 [5] S. N. Gninenko, *Phys. Rev. D* **83**, 015015 (2011).

[6] D. McKeen and M. Pospelov, *Phys. Rev. D* **82**, 113018 (2010).
 [7] M. Masip, P. Masjuan, and D. Meloni, *J. High Energy Phys.* **01** (2013) 106.
 [8] M. Masip and P. Masjuan, *Phys. Rev. D* **83**, 091301 (2011).
 [9] S. N. Gninenko, *Phys. Lett. B* **710**, 86 (2012).

- [10] C. Giunti and A. Studenikin, *Rev. Mod. Phys.* **87**, 531 (2015).
- [11] A. Aparici, K. Kim, A. Santamaria, and J. Wudka, *Phys. Rev. D* **80**, 013010 (2009).
- [12] A. Aparici, Exotic properties of neutrinos using effective Lagrangians and specific models, Ph.D. thesis, Valencia U., 2013.
- [13] A. Caputo, P. Hernandez, J. Lopez-Pavon, and J. Salvado, *J. High Energy Phys.* **06** (2017) 112.
- [14] P. Coloma, P. A. N. Machado, I. Martinez-Soler, and I. M. Shoemaker, *Phys. Rev. Lett.* **119**, 201804 (2017).
- [15] K. N. Abazajian, *Phys. Rep.* **711–712**, 1 (2017).
- [16] C. Athanassopoulos *et al.* (LSND Collaboration), *Phys. Rev. Lett.* **77**, 3082 (1996).
- [17] A. A. Aguilar-Arevalo *et al.* (MiniBooNE Collaboration), *Phys. Rev. Lett.* **98**, 231801 (2007).
- [18] R. J. Hill, *Phys. Rev. D* **84**, 017501 (2011).
- [19] K. N. Abazajian *et al.*, arXiv:1204.5379.
- [20] F. Couchot, S. Henrot-Versillé, O. Perdereau, S. Plaszczynski, B. R. D’Orfeuille, M. Spinelli, and M. Tristram, *Astron. Astrophys.* **606**, A104 (2017).
- [21] B. Dasgupta and J. Kopp, *Phys. Rev. Lett.* **112**, 031803 (2014).
- [22] S. Hannestad, R. S. Hansen, and T. Tram, *Phys. Rev. Lett.* **112**, 031802 (2014).
- [23] M. Antonello *et al.* (LAr1-ND, ICARUS-WA104, and MicroBooNE Collaborations), arXiv:1503.01520.
- [24] P. Ballett, S. Pascoli, and M. Ross-Lonergan, *J. High Energy Phys.* **04** (2017) 102.
- [25] V. Barger, W.-Y. Keung, and D. Marfatia, *Phys. Lett. B* **696**, 74 (2011).
- [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, *Nucl. Phys.* **B206**, 359 (1982).
- [29] R. N. Mohapatra, *Phys. Rev. Lett.* **56**, 561 (1986).
- [30] R. N. Mohapatra and J. W. F. Valle, *Phys. Rev. D* **34**, 1642 (1986).
- [31] B. Grzadkowski, M. Iskrzynski, M. Misiak, and J. Rosiek, *J. High Energy Phys.* **10** (2010) 085.
- [32] A. A. Aguilar-Arevalo *et al.* (MiniBooNE Collaboration), *Phys. Rev. D* **79**, 072002 (2009).
- [33] L. B. Auerbach *et al.* (LSND Collaboration), *Phys. Rev. D* **63**, 112001 (2001).
- [34] E. Izaguirre, G. Krnjaic, and M. Pospelov, *Phys. Lett. B* **740**, 61 (2015).
- [35] Y. Kahn, G. Krnjaic, J. Thaler, and M. Tups, *Phys. Rev. D* **91**, 055006 (2015).
- [36] M. Pospelov and Y.-D. Tsai, *Phys. Lett. B* **785**, 288 (2018).
- [37] R. Burman, M. Potter, and E. Smith, *Nucl. Instrum. Methods Phys. Res., Sect. A* **291**, 621 (1990).
- [38] R. Burman and E. Smith, LAMPF Report No. LA-11502-MS, 1989.
- [39] V. Bernard, T. R. Hemmert, and U.-G. Meissner, *Nucl. Phys.* **A686**, 290 (2001).
- [40] L. Alvarez-Ruso and E. Saul-Sala, in *Proceedings of the Prospects in Neutrino Physics (NuPhys2016), London, UK, 2016* (2017) [arXiv:1705.00353].
- [41] F. Vannucci, *Adv. High Energy Phys.* **2014**, 129694 (2014).
- [42] J. Altegoer *et al.*, *Nucl. Instrum. Methods Phys. Res., Sect. A* **404**, 96 (1998).
- [43] J. Altegoer *et al.*, *Phys. Lett. B* **428**, 197 (1998).
- [44] S. N. Gninenko and N. V. Krasnikov, *Phys. Lett. B* **450**, 165 (1999).
- [45] S. Gninenko and N. Krasnikov, *Phys. Lett. B* **427**, 307 (1998).
- [46] M. Anelli *et al.* (SHiP Collaboration), arXiv:1504.04956.
- [47] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, *J. High Energy Phys.* **07** (2014) 079.
- [48] A. Alloul, N. D. Christensen, C. Degrande, C. Duhr, and B. Fuks, *Comput. Phys. Commun.* **185**, 2250 (2014).
- [49] N. D. Christensen and C. Duhr, *Comput. Phys. Commun.* **180**, 1614 (2009).
- [50] O. Adriani *et al.*, *Phys. Lett. B* **297**, 469 (1992).
- [51] R. Akers *et al.* (OPAL Collaboration), *Z. Phys. C* **65**, 47 (1995).
- [52] P. Abreu *et al.* (DELPHI Collaboration), *Z. Phys. C* **74**, 577 (1997).
- [53] J. L. Lopez, D. V. Nanopoulos, and A. Zichichi, *Phys. Rev. D* **55**, 5813 (1997).
- [54] R. Assmann, M. Lamont, and S. Myers, *Nucl. Phys. B, Proc. Suppl.* **109**, 17 (2002).
- [55] M. Aaboud *et al.* (ATLAS Collaboration), *Eur. Phys. J. C* **77**, 393 (2017).
- [56] M. Aaboud *et al.* (ATLAS Collaboration), *Eur. Phys. J. C* **76**, 666 (2016).
- [57] V. Khachatryan *et al.* (CMS Collaboration), *Phys. Lett. B* **757**, 6 (2016).
- [58] O. Ruchayskiy and A. Ivashko, *J. Cosmol. Astropart. Phys.* **10** (2012) 014.
- [59] S. Dodelson and L. M. Widrow, *Phys. Rev. Lett.* **72**, 17 (1994).
- [60] B. D. Fields, P. Molaro, and S. Sarkar, *Chin. Phys. C* **38**, 339 (2014).
- [61] G. G. Raffelt, *Stars as Laboratories for Fundamental Physics* (University of Chicago Press, Chicago, 1996).
- [62] H. K. Dreiner, C. Hanhart, U. Langenfeld, and D. R. Phillips, *Phys. Rev. D* **68**, 055004 (2003).
- [63] H. K. Dreiner, H. E. Haber, and S. P. Martin, *Phys. Rep.* **494**, 1 (2010).
- [64] H. K. Dreiner, J.-F. Fortin, C. Hanhart, and L. Ubaldi, *Phys. Rev. D* **89**, 105015 (2014).
- [65] T. Fischer, S. Chakraborty, M. Giannotti, A. Mirizzi, A. Payez, and A. Ringwald, *Phys. Rev. D* **94**, 085012 (2016).
- [66] J. H. Chang, R. Essig, and S. D. McDermott, *J. High Energy Phys.* **01** (2017) 107.
- [67] E. Hardy and R. Lasenby, *J. High Energy Phys.* **02** (2017) 033.
- [68] K. Hirata *et al.*, *Phys. Rev. Lett.* **58**, 1490 (1987).
- [69] E. N. Alekseev, L. N. Alekseeva, I. V. Krivosheina, and V. I. Volchenko, *Phys. Lett. B* **205**, 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. D* **89**, 103533 (2014).
- [73] H. K. Dreiner, C. Hanhart, U. Langenfeld, and D. R. Phillips, *Phys. Rev. D* **68**, 055004 (2003).
- [74] A. Fradette and M. Pospelov, *Phys. Rev. D* **96**, 075033 (2017).
- [75] S. A. Colgate and R. H. White, *Astrophys. J.* **143**, 626 (1966).

- [76] M. Rapp and H. T. Janka, *Astrophys. J.* **539**, L33 (2000).
- [77] M. Liebendoerfer, A. Mezzacappa, F.-K. Thielemann, O. E. B. Messer, W. R. Hix, and S. W. Bruenn, *Phys. Rev. D* **63**, 103004 (2001).
- [78] B. Müller, T. Melson, A. Heger, and H. T. Janka, *Mon. Not. R. Astron. Soc.* **472**, 491 (2017).
- [79] G. T. Zatsepin and A. Yu. Smirnov, *Pis'ma Zh. Eksp. Teor. Fiz.* **28**, 379 (1978); [*J. Exp. Theor. Phys.* **28**, 379 (1978)].
- [80] G. M. Fuller, A. Kusenko, and K. Petraki, *Phys. Lett. B* **670**, 281 (2009).
- [81] J. P. Chou, D. Curtin, and H. J. Lubatti, *Phys. Lett. B* **767**, 29 (2017).
- [82] V. V. Gligorov, S. Knapen, M. Papucci, and D. J. Robinson, *Phys. Rev. D* **97**, 015023 (2018).
- [83] J. Feng, I. Galon, F. Kling, and S. Trojanowski, *Phys. Rev. D* **97**, 035001 (2018).
- [84] F. Kling and S. Trojanowski, *Phys. Rev. D* **97**, 095016 (2018).
- [85] J. C. Helo, M. Hirsch, and Z. S. Wang, *J. High Energy Phys.* **07** (2018) 056.
- [86] G. Magill, R. Plestid, M. Pospelov, and Y.-D. Tsai, [arXiv:1806.03310](https://arxiv.org/abs/1806.03310) [*Phys. Rev. Lett.* (to be published)].
- [87] G. Magill and R. Plestid, *Phys. Rev. D* **95**, 073004 (2017).
- [88] U. D. Jentschura and V. G. Serbo, *Eur. Phys. J. C* **64**, 309 (2009).
- [89] C. F. Perdrisat, V. Punjabi, and M. Vanderhaeghen, *Prog. Part. Nucl. Phys.* **59**, 694 (2007).
- [90] D. H. Beck and B. R. Holstein, *Int. J. Mod. Phys. E* **10**, 1 (2001).
- [91] P. deNiverville, C.-Y. Chen, M. Pospelov, and A. Ritz, *Phys. Rev. D* **95**, 035006 (2017).
- [92] M. Bonesini, A. Marchionni, F. Pietropaolo, and T. T. de Fatis, *Eur. Phys. J. C* **20**, 13 (2001).
- [93] C. Mariani, G. Cheng, J. M. Conrad, and M. H. Shaevitz, *Phys. Rev. D* **84**, 114021 (2011).
- [94] D. A. Dicus and W. W. Repko, *Phys. Rev. Lett.* **79**, 569 (1997).
- [95] A. Abbasabadi, A. Devoto, and W. W. Repko, *Phys. Rev. D* **63**, 093001 (2001).
- [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. Rep.* **15**, 181 (1975).
- [98] G. Magill and R. Plestid, *Phys. Rev. D* **97**, 055003 (2018).
- [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), *Phys. Lett. B* **667**, 1 (2008). K. Nakamura *et al.* (Particle Data Group), *J. Phys. G* **37**, 075021 (2010).
- [102] L. F. Li and F. Wilczek, *Phys. Rev. D* **25**, 143 (1982).
- [103] C. Patrignani *et al.* (Particle Data Group), *Chin. Phys. C* **40**, 100001 (2016).