

Model agnostic probes of dark sectors at neutrino experimentsMarco Costa^{1,2}, Rashmish K. Mishra³ and Sonali Verma^{1,2}¹*Scuola Normale Superiore, Piazza dei Cavalieri 7, 56126 Pisa, Italy*²*INFN Sezione di Pisa, Largo Bruno Pontecorvo 3, 56127 Pisa, Italy*³*Harvard University, 17 Oxford Street, Cambridge, Massachusetts 02138, USA*

(Received 10 April 2023; accepted 27 July 2023; published 25 August 2023)

Present and upcoming neutrino experiments can have considerable sensitivity to dark sectors that interact feebly with the Standard Model (SM). We consider dark sectors interacting with the SM through irrelevant portals that are motivated on general principles. We derive bounds on such scenarios by considering decays of dark sector excitations inside the neutrino detector, placed downstream from the target. Our approach is model agnostic and applies to a wide range of dark sector models, both strongly and weakly coupled. In this approach, the dark sector is characterized by two scales; Λ_{UV} (mass of mediators generating the portals) and Λ_{IR} (mass gap of the dark sector). At intermediate energies, far away from these scales, the theory is approximately scale invariant. This allows the calculation of production rates independent of the threshold corrections, although some mild model-dependent assumptions are needed. We look at various dark sector production processes relevant at neutrino experiments such as meson decays, direct partonic production, and dark bremsstrahlung. We consider representative experiments from past (CHARM), present (ICARUS, NOvA, MicroBooNE), and upcoming future (DUNE-MPD), and compare their reach to existing bounds from high-energy experiments (LHC and LEP) and dedicated future LLP experiments (SHiP). We find that the upcoming DUNE-MPD can probe Λ_{UV} in the TeV range, and Λ_{IR} in the 0.1–1 GeV range, covering parts of parameter space currently inaccessible in high-energy experiments and fixed-target/beam dump experiments, and is comparable to future LLP experiments. In general, future neutrino experiments can be an efficient probe of dark sectors, providing complementary as well as a new reach in parameter space.

DOI: [10.1103/PhysRevD.108.035041](https://doi.org/10.1103/PhysRevD.108.035041)**I. INTRODUCTION**

Secluded sectors with their own particle content and dynamics, that interact with the Standard Model (SM) feebly, are well-motivated for a variety of reasons. If such sectors contain or interact with dark matter, they are a natural scenario to probe in terrestrial and cosmological studies. Such sectors can also arise naturally from bottom-up beyond the standard model (BSM) considerations motivated to address various issues in SM, as well as in top-down string constructions. The hidden valley scenario [1] initiated various model-building and phenomenological aspects of such dark sectors (DS), and this remains a focus of much present research activity to date.

Dark sectors with sufficiently low production threshold are generally probed by high-intensity experiments with center-of-mass energies much lower than typical high-energy collider machines. Such high-intensity experiments

involve a proton or electron beam (or even muon beam, see Refs. [2–4] for future proposals of fixed-target experiments with a muon beam) hitting a fixed target, producing a high flux of SM particles alongside a beam of putative DS particles (see Refs. [5–10] for general reviews on the capabilities of high-intensity experiments). Due to weak couplings to the SM, lightest DS particles once produced tend to have long lifetimes, allowing them to travel macroscopic distances before decaying back to visible particles. Suitable detectors placed downstream from the fixed target can be used to discriminate a possible DS signal against SM background.

Short- and long-baseline neutrino experiments happen to be placed behind some of the most powerful proton beams up to date. Thus, they provide an ideal and preexisting infrastructure for probing low-scale dark sectors. In fact, neutrinos themselves are a prototype for a DS, so it is not a surprise that a facility for studying them is useful more generally. Further, a rich short-baseline experimental program for neutrinos has been planned at Fermilab (e.g., DUNE, SBNP, see Refs. [11–13] for experiment details, Refs. [14–16] for recent studies on DS search at these experiments). These new proposals will improve upon the current neutrino experiments, using a higher number of

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. 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.

protons on target (POT), and better detectors that can help in reducing SM background [17], therefore leading to an increased sensitivity for DS searches.

Apart from neutrino experiments, a natural setup to probe DS is at future experiments proposed for long-lived particle (LLP) searches (e.g., see Refs. [18,19] for a review). These experiments are in the approval phase and may require a longer timescale. Existence of current data from neutrino experiments, and a fairly short timescale for future ones to come online make neutrino experiments a powerful and efficient probe of dark sectors.

Any broad-enough search for dark sectors must explore all axes of ignorance of such scenarios. The dimensionality D of the portal interaction between the DS and SM is one such axis. While searches for DS interacting with the SM via relevant portals ($D \leq 4$) have been well-studied (see Refs. [19–21] and references therein), the case of irrelevant portals ($D > 4$) is equally well-motivated. Among these, the axion portal ($D = 5$) has been studied the most, especially at the high-intensity frontier [22–28] (see Ref. [29] for a study specific to neutrino experiments). Regarding other irrelevant portals ($D \geq 5$) that might connect the DS and SM, some recent progress has been made [30–37]. In this work, we probe SM-neutral dark sectors which interact via irrelevant portals with the SM at neutrino-oscillation experiments. On general principles we are led to $D \sim 5$ and $D \sim 6$ portals. Such dark sectors are in general very elusive due to the irrelevant nature of the portal. The results presented here are complementary to the constraints from current high-energy terrestrial experiments and astrophysical data, presented in [30], and in specific cases, much stronger, as we point out in the relevant sections.

In this work we consider various (inclusive) DS production processes. Depending on the DS 4-momentum p_{DS} , a different production process can be relevant; meson decays (for $\sqrt{p_{\text{DS}}^2} \lesssim M$, the parent meson mass), direct partonic production (for $\sqrt{p_{\text{DS}}^2} \gtrsim \Lambda_{\text{QCD}}$), or dark bremsstrahlung (for $\sqrt{p_{\text{DS}}^2} < \Lambda_{\text{QCD}}$). We require the DS states to be produced away from any mass thresholds, which allows estimating the rates based on general principles. DS particles once produced are required to decay to SM particles inside the neutrino near-detectors placed generally $\sim \mathcal{O}(100 \text{ m})$ downstream from the target.¹

Compared to previous attempts at studying irrelevant portals, our work is more comprehensive, as we point out now. Compared to [26,33], we study a more complete set of operators, both in production and decay of DS particles. We perform a detailed study of production modes through irrelevant portals, such as dark bremsstrahlung and partonic production, that have either been neglected or only considered partially [33–35]. We find that the bremsstrahlung

mode can be comparable to other modes and is necessary for a complete analysis. Compared to [30,32] we focus on high-intensity experiments, particularly on neutrino- and other proton-dump experiments that were not considered previously in this framework. We do so by adopting the model-agnostic strategy outlined in [30]. More importantly, this allows us to put bounds on strongly coupled light-dark sectors through irrelevant portals, which as far as we know is not a thoroughly studied scenario at high-intensity experiments (except in [30]. See Refs. [32,36,38,39] and references therein for searches at collider experiments). Strongly coupled GeV-scale DS are relevant in frameworks containing composite resonances from a new gauge group, such as composite versions of asymmetric dark matter [40–42], mirror world models [43,44], some incarnations of the twin Higgs paradigm [45], and are a natural realization of the hidden valley scenario [1].

The outline of the paper is as follows. In Sec. II we describe the dark sector portals and the relevant model-agnostic framework for estimating inclusive rates, lifetime and multiplicity of DS particles, also pointing out the mild model-dependent assumptions we have to make to proceed. In Sec. III we describe various DS production processes relevant at neutrino experiments, while in Sec. IV, we give an overview of the neutrino experiments we use to constrain the parameter space, and describe our strategy for estimating signal events from DS decaying inside the neutrino detector. Our results and bounds can be found in Sec. V with a discussion and summary in Sec. VI. Appendixes A–E contain technical details.

II. A MODEL AGNOSTIC STRATEGY (AND ITS LIMITATIONS)

In this section we discuss the relevant theoretical details for studying dark sectors with irrelevant portal to the SM. The emphasis is towards being as model agnostic as possible, and only allowing for minimal model dependence where necessary. We point out the assumptions we have to make at various stages for this. Our work builds upon the model agnostic approach first undertaken in [30], wherein more details can be found.

Dark sectors with portal interactions to the SM from irrelevant operators can be generated in a large class of models, generically by exchange of heavy mediators charged under both the SM and the DS. The general form of such a portal is

$$\frac{\kappa}{\Lambda_{\text{UV}}^{D-4}} \mathcal{O}_{\text{SM}} \mathcal{O}_{\text{DS}}, \quad (1)$$

where Λ_{UV} is the mass scale of the heavy mediator, κ is a dimensionless coupling, and \mathcal{O}_{SM} (\mathcal{O}_{DS}) are local operators made of SM (DS) degrees of freedom. The dimensionality of portal $D = [\mathcal{O}_{\text{SM}}] + [\mathcal{O}_{\text{DS}}]$ is greater than 4 for irrelevant portals. States in the DS are further characterized by a mass

¹Far detectors are less constraining due to a very small angle subtended to the interaction point.

gap Λ_{IR} , and the dynamics between the scales Λ_{UV} and Λ_{IR} is approximately scale invariant. A large hierarchy between these scales is a working assumption of this scenario. In order to avoid strong constraints, we also assume that the portal preserves both CP and flavor symmetries of the SM.

The most constraining portals are expected to be those with lowest dimension D . In concrete examples of such DS, both weakly and strongly coupled, the two lowest-dimension DS operators are a scalar operator \mathcal{O} (of dimension $\Delta_{\mathcal{O}} \lesssim 4$) and a conserved current operator J_{μ}^{DS} (of dimension 3). While the dimension of a conserved current operator is fixed to be 3 in $4D$, the reason to take the scalar slightly marginal is to ensure that the condition $\Lambda_{\text{IR}}/\Lambda_{\text{UV}} \ll 1$ is realized naturally. We will consider $\Delta_{\mathcal{O}} = 3, 4$ in this work. Specific to a given model there can be other operators that generate portals to the SM. However in the absence of a symmetry, their dimension is either unprotected or requires additional assumptions about the DS. We will therefore limit ourselves with only a current and a scalar operator on the DS side. The gauge-invariant operators on the SM side that can be used to make a portal operator, with increasing scaling dimension, are $H^{\dagger}H$ and $J_{\mu}^{\text{SM}} = \bar{f}\gamma_{\mu}f, \bar{f}\gamma_{\mu}\gamma^5 f, f = l, q$ or $J_{\mu}^{\text{SM}} = H^{\dagger}i\overleftrightarrow{D}_{\mu}H$ (see Table 1 in Ref. [30] for a complete list of scenarios). The lowest-dimension Lorentz-invariant combinations are then $J_{\mu}^{\text{DS}}J_{\text{SM}}^{\mu}$, and $\mathcal{O}H^{\dagger}H$. In the unitary gauge, $H^{\dagger}i\overleftrightarrow{D}_{\mu}H \sim Z_{\mu}$ so that for this portal the interactions with the DS proceed through a Z -boson. We will refer to this as the Z portal. At the energy scales relevant at neutrino experiments, Z is never produced on shell, so we can integrate it out and generate an effective JJ operator where now the SM current is the one that couples to Z . Therefore, considering JJ portals where the SM current is either generic or the one for Z , we cover all possibilities. We will refer to these as the generic JJ portal and the Z -aligned JJ portal respectively.

Hence the lowest-dimension portals that can be formed are

$$\begin{aligned} \mathcal{L}_{\text{portal}} &= \frac{\kappa_{\mathcal{O}}}{\Lambda_{\text{UV}}^{\Delta_{\mathcal{O}}-2}} \mathcal{O}H^{\dagger}H + \frac{\kappa_J}{\Lambda_{\text{UV}}^2} J_{\mu}^{\text{DS}}J_{\text{SM}}^{\mu} + \frac{\kappa_Z}{\Lambda_{\text{UV}}^2} J_{\mu}^{\text{DS}}H^{\dagger}i\overleftrightarrow{D}^{\mu}H \\ &= \frac{\kappa_{\mathcal{O}}}{\Lambda_{\text{UV}}^{\Delta_{\mathcal{O}}-2}} \mathcal{O}H^{\dagger}H + \frac{\kappa_J}{\Lambda_{\text{UV}}^2} J_{\mu}^{\text{DS}}J_{\text{SM}}^{\mu} + \frac{\kappa_Z}{\Lambda_{\text{UV}}^2} \frac{v_{\text{EW}}}{m_Z} J_{\mu}^{\text{DS}}J_{\text{SM},Z}^{\mu}, \end{aligned} \quad (2)$$

where $\kappa_{\mathcal{O}}, \kappa_J, \kappa_Z$ are dimensionless coefficients, v_{EW} is the electroweak VEV and in the second line we have integrated out Z , which couples the DS current to $J_{\text{SM},Z}^{\mu}$, the SM current that couples to Z . The three terms in Eq. (2) are the Higgs portal, the generic JJ portal, and the Z portal, respectively. It is clear that a Z -aligned JJ portal can be obtained from the Z -portal with a rescaling; $\kappa_J = \kappa_Z(v_{\text{EW}}/m_Z)$. For $\Delta \lesssim 4$, all these portals are of dimension $D \sim 6$. In principle, a DS described by a local quantum field theory (QFT) also

possesses a stress-energy tensor $T_{\text{DS}}^{\mu\nu}$ of dimension 4, that can be used to build dimension-8 operators with SM dimension-4 operators. However, given a larger suppression compared to the dimension-6 portals in (2), the bounds on them are too weak to be of any interest.

If the energy \sqrt{s} of an experiment that probes the DS is such that $\Lambda_{\text{IR}} \ll \sqrt{s} \ll \Lambda_{\text{UV}}$, the DS states are produced directly in the conformal regime. Inclusive DS production rates can be estimated using only the scaling dimension of the DS operator, along with the optical theorem. The optical theorem allows to sum over the DS phase space in an inclusive manner and relates it to the imaginary part of the two point function of the DS operator, which in turn is fixed by the scaling dimension of the operator. In particular, the optical theorem gives

$$\sum_n \int d\Phi_{\text{DS}} |\langle \Omega | \mathcal{O}_{\text{DS}} | n \rangle|^2 = 2\text{Im}(i\langle \Omega | \mathcal{O}_{\text{DS}} \mathcal{O}_{\text{DS}} | \Omega \rangle), \quad (3)$$

where the DS operator \mathcal{O}_{DS} interpolates a DS state $|n\rangle$ from vacuum $|\Omega\rangle$ and the integration is over the entire dark-sector phase space $d\Phi_{\text{DS}}$.

This approach allows calculating the cross section for DS production without specifying the fields that make the composite operator \mathcal{O}_{DS} . While for irrelevant portals, the matrix element does not decrease with energy (specific behavior depends on the production mode), this needs to be convolved with the structure functions (e.g., the parton distribution functions/form factors/splitting functions, depending on the production channel), and this changes where the bulk of events come from. As long as the involved p_{DS} values are away from $\Lambda_{\text{UV}}, \Lambda_{\text{IR}}$, one can ignore the events near the thresholds in a self-consistent manner. Relatedly, the two-point function of \mathcal{O}_{DS} will also depend on the ratio $\Lambda_{\text{IR}}^2/p_{\text{DS}}^2$ and $p_{\text{DS}}^2/\Lambda_{\text{UV}}^2$. For self-consistency, we again need both these ratios to be small. In particular, the condition $p_{\text{DS}}^2/\Lambda_{\text{UV}}^2 < 1$ effectively ensures the mediators of mass Λ_{UV} are not directly produced and the effective local operator for the portal is a good description. In Ref. [30], this was enforced by ensuring that the obtained bound on Λ_{UV} always satisfies this condition for the highest p_{DS}^2 used in the calculation. In practice, this effectively resulted in a lower limit on the parts of Λ_{UV} ruled out, or completely invalidated certain bounds. As we will see, for neutrino experiments, where the involved energy is much smaller than LHC or LEP, this condition is less detrimental. By restricting to $\Lambda_{\text{UV}} \gtrsim 50$ GeV, we are able to get useful bounds as well as be consistent with the effective field theory (EFT) condition. The condition on Λ_{IR} on the other hand needs to be imposed, which we do for each production mode.

After production, the DS states will interact and decay among each other, and eventually all the DS degrees of freedom would decay to the lightest dark sector particle (LDSP), which we denote by ψ . We will take the mass of ψ to be of order Λ_{IR} and this can be taken as our definition for

the mass gap Λ_{IR} . In the absence of additional symmetry, ψ will decay back to SM states from the portal interactions itself. Since the portal interactions are weak, the typical time for DS states to decay among each other is much smaller than the typical lifetime of ψ , and can be safely ignored. Note that the LDSP is not the DM candidate in the scenario under consideration—a DM candidate would need to be much more long-lived, and will have a missing energy signal. In this work we will assume that the DS relaxes entirely to LDSPs, and leave the question of considering a fraction of events to be missing energy, for future work.

The signatures of ψ depend on its lifetime, and this is the first place where some assumptions have to be made, which bring some model dependence. At high-energy colliders, depending on the lifetime of ψ , one can get missing-energy events, displaced vertices, or prompt decays, ordered by decreasing lifetimes. Missing-energy events, being most inclusive, need minimal information about the underlying dynamics of DS, while displaced vertices and prompt decays being exclusive, need some information. Note that the requirement $\Lambda_{\text{IR}}/\Lambda_{\text{UV}} \ll 1$ puts us away from the prompt-decay regime, since in this limit, the lifetime increases. Focusing on neutrino experiments, since the detectors are placed some distance from the interaction point, we are in the displaced vertex scenario. It is possible to detect the decay of ψ inside the detector, or its scattering against electrons or nucleons of the detector [46–53]. Both signatures need some knowledge of the IR behavior of the underlying theory, and are model dependent; however, with varying degrees. For these two signatures, the relevant DS matrix elements are

$$\text{Decay: } \langle \Omega | \mathcal{O}_{\text{DS}} | \psi \rangle = af \Lambda_{\text{IR}}^{\Delta_{\mathcal{O}} - 2},$$

$$\text{Scattering: } \langle \psi(p_f) | \mathcal{O}_{\text{DS}} | \psi(p_i) \rangle = aF(p_i, p_f, \Lambda_{\text{IR}}), \quad (4)$$

where $|\Omega\rangle$ is the vacuum, f is a decay constant, F is a form factor, and a is an $\mathcal{O}(1)$ number, all of which are model dependent. In this work, we will only focus on the decay mode, as we explain this choice now.

The model dependence that enters in the decay case comes in the combination af . The scattering process, on the other hand, requires knowing the form factor which can be a complicated function of the momenta (especially for strongly coupled sectors). The functional dependence also influences who ψ recoils against most efficiently. The spin of ψ does not fix the portal, since one can make multiple total spin states using two ψ . Further, depending on the spectrum, an LDSP might up scatter to a close by state, making the scattering inelastic (similar to what happens in inelastic dark matter scenarios [54]), leading to a different parametric dependence for the scattering cross section. These aspects make it clear that scattering processes require additional model-dependent assumptions, and we will not consider them here. A further reason to choose decays over scattering is that they have a larger signal-to-noise ratio,

and we will have more to say about it in Sec. IV. Note that while there are weakly coupled models in which all the LDSPs are stable under some accidental symmetry, and therefore can only be studied through scatterings in the experiments under scrutiny (and therefore our analysis will not apply to such scenarios), in strongly coupled models unstable resonances are expected generically.

The lifetime of ψ to decay to SM states, via the portal itself, can be estimated in a straightforward manner. However there are differences when the decay is from mixing with a SM state or a direct decay. For a direct decay from a portal of dimension D , the lifetime can be estimated to be

$$\frac{1}{\tau_\psi} \sim \Lambda_{\text{IR}} \frac{\kappa^2 f^2}{8\pi \Lambda_{\text{IR}}^2} \left(\frac{\Lambda_{\text{IR}}^2}{\Lambda_{\text{UV}}^2} \right)^{D-4}, \quad (5)$$

where the decay constant f is defined by the matrix element $\langle \Omega | \mathcal{O} | \psi \rangle = af \Lambda_{\text{IR}}^{\Delta_{\mathcal{O}} - 2}$ and a is an $\mathcal{O}(1)$ number taken to be 1. Further, f can be estimated to be $f = \sqrt{c} \Lambda_{\text{IR}} / 4\pi$, where c is the number of degrees of freedom of the DS. On the other hand, if the LDSP decays through mixing with a SM particle such as Higgs, if the LDSP is spin 0, or Z , if the LDSP is spin 1 (e.g., through $\mathcal{O}H^\dagger H$ or $J_\mu^{\text{DS}} H^\dagger i \overleftrightarrow{D}^\mu H$ respectively), the lifetime in the limit $\Lambda_{\text{IR}} \ll m_{Z/h}$ is given as

$$\frac{1}{\tau_\psi} = \Gamma_i \sin^2 \theta_i, \quad \tan 2\theta_i = \frac{2\delta_i}{m_i^2}, \quad i = Z, h, \quad (6)$$

where $\Gamma_{Z/h}$ is the decay width of Z/h evaluated at $m_{Z/h} = \Lambda_{\text{IR}}$, and the mixing parameter δ_i is

$$\begin{aligned} \delta_h &= \kappa_{\mathcal{O}} v_{\text{EW}} f \left(\frac{\Lambda_{\text{IR}}}{\Lambda_{\text{UV}}} \right)^{\Delta_{\mathcal{O}} - 2}, \\ \delta_Z &= \kappa_J v_{\text{EW}} f \frac{m_Z \Lambda_{\text{IR}}}{\Lambda_{\text{UV}}^2}. \end{aligned} \quad (7)$$

To model the hadronic decay of the scalar LDSP (that mixes with the Higgs), we use the spectator quark model for $\Lambda_{\text{IR}} > 2$ GeV and the dispersive analysis for $\Lambda_{\text{IR}} < 2$ GeV, following [55]. For a spin-1 LDSP (mixing with the Z) we again use the spectator quark model for $\Lambda_{\text{IR}} > 2$ GeV, and a data-driven approach for $\Lambda_{\text{IR}} < 2$ GeV, following [56,57] for the vector and axial vector component respectively.

The next model dependent assumption needed in order to evaluate the reach at high intensity experiments is how many LDSPs are produced per DS shell, or equivalently how many are excited by the DS operator acting on the vacuum. We will take two benchmark values, $n_{\text{LDSP}} = 2$ for weakly coupled dark sectors and $n_{\text{LDSP}} = n(p_{\text{DS}}^2)$ a function of the invariant mass squared p_{DS}^2 of the DS system, similar to the case of QCD [58],

$$n(p_{\text{DS}}^2) = A (\log x)^B \exp\{C(\log x)^D\}, \quad (8)$$

where $x = p_{\text{DS}}^2/\bar{\Lambda}^2$, $\bar{\Lambda} = 0.1\Lambda_{\text{IR}}$, $A = 0.06$, $B = -0.5$, $C = 1.8$, $D = 0.5$, and p_{DS}^2 is the invariant mass squared of the DS system. Our results are not very sensitive to small changes in n_{LDSP} . In particular, as argued in Sec. V, its impact on the exclusion plots will be mostly in regions in which the LDSPs are light and long-lived.

Finally, we need to know the directional distribution of the produced LDSPs, to estimate if they interact with the detector. In the strongly coupled benchmark where typically $n_{\text{LDSP}} > 2$, we assume that the LDSPs have a uniform angular distribution in the rest frame of DS (i.e., the frame in which p_{DS} only has a time component), and we can boost it to the lab frame to know its relevant distribution. A uniform distribution in the rest frame is a simplifying choice, and is well-motivated, at least for a certain class of strongly coupled theories (e.g., see Ref. [59] for such a scenario). Further, even if the distribution is not uniform per event, it can be uniform when all the events are considered. For light enough LDSPs, which are very boosted in the lab frame, small deviations from this assumption do not change our results significantly.

The weakly coupled case is in principle different, and the angular distribution depends on the production mode, spin of produced DS particles, and the specific form of the portal. In general we expect $\mathcal{O}(1)$ differences among the possible LDSP angular distributions in the DS rest frame. For example, in Drell-Yan (DY) production the typical LDSP distribution is either proportional to $\sin^2\theta$ or $1 + \cos^2\theta$ for scalar and light fermion LDSPs respectively. The difference between the two distributions is that the scalar distribution is more peaked around the most probable LDSP lab angle $\sim 1/\gamma_{\text{DS}}$. However, since the LDSP is produced with a high boost, any differences in the distribution are washed out, and we can assume an isotropic distribution in the DS rest frame as before. We have checked this by an explicit computation.

III. PRODUCTION MODES AT PROTON BEAM-BASED EXPERIMENTS

Even though at neutrino experiments the primary process is a proton interacting with a nucleus, depending on the energy scale of the process, there are various production modes to consider. In this work, we consider experiments based on 120 GeV and 400 GeV beam energies. For such energies there are three relevant production modes. First of all, the proton nucleus interaction creates mesons, which may decay into lighter mesons and DS states, or completely annihilate into DS states. Denoting the 4-momentum carried by the DS state as p_{DS} , this requires $p_{\text{DS}}^2 \leq (\mathbf{M}_{\text{heavy}} - \mathbf{M}_{\text{light}})^2$ for the first scenario and $p_{\text{DS}}^2 = \mathbf{M}_{\text{heavy}}^2$ for the second. We will refer to these as radiative and annihilation meson decays (MD)

respectively. For $p_{\text{DS}}^2 \gtrsim \Lambda_{\text{QCD}}^2$, the incoming proton is at high enough energies so one has to consider the partonic process involving constituents from the incoming proton and the nucleons in the target, and we refer to this as the Drell-Yan production mode. For $p_{\text{DS}}^2 \lesssim \Lambda_{\text{QCD}}^2$, DS states can be produced from initial state emission, which we will refer to as dark bremsstrahlung (DB) mode. For each of these processes, the production cross section has a different differential distribution in p_{DS}^2 . Figure 1 shows a comparison of the differential DS production cross section for DY, DB, and radiative MD mode, for Z portal, at 120 GeV beam energy. The radiative MD mode is flat in p_{DS} , switching off when the phase space for DS production closes, which in turn is set by the parent meson mass. The DB mode switches off around Λ_{QCD} beyond which it is not a valid description of the scattering process. The sharp peak in the DB mode is due to meson resonance, as seen in the form factors (see Appendix A). The switch off of DY mode comes from the drop in the parton distribution functions (PDFs) of constituents of the proton at higher p_{DS}^2 , given \sqrt{s} of the experiment, and is a slower drop.

We remark that dependence of the production cross section on the center of mass energy \sqrt{s} of the experiment is not the one given by naive power counting (i.e., $\sigma \sim s^{D-5}/\Lambda_{\text{UV}}^{2D-8}$), and is general dependent on the production mode. For MD and DB modes, the typical scale of the process is not set by the center-of-mass energy of the experiment, but rather by the hadronic resonances. There is a residual dependence on \sqrt{s} in the meson production cross section and in the proton-nucleon cross section respectively, but typically this dependence is much weaker than




FIG. 1. Relative importance of various production modes: the scaled differential cross section for DS production at DUNE-MPD ($E_{\text{beam}} = 120$ GeV) as a function of p_{DS}^2 for various DS production modes (for Z portal). Solid yellow shows meson decay $K \rightarrow \pi + \text{DS}$, dotted green line shows dark bremsstrahlung mode ($p + p \rightarrow X + \text{DS}$), and dashed blue line shows Drell-Yan mode ($p + p \rightarrow \text{DS}$). The reported cross-section is per proton-on-target, and is without the geometric acceptance factor ϵ_{geom} (which at DUNE is approximately 10^{-3} for DY and meson modes and around 10^{-2} for DB mode).

the expected one above the specific hadronic production threshold. DY production instead is more sensitive on \sqrt{s} . These features are easily seen in Fig. 1.

Independent of the production mode, we need to estimate the number of DS signal events S produced. This is generically given as

$$S = N_{\text{signal}} = N_{\text{DS}} P_{\text{decay}} \epsilon_{\text{geo}}, \quad (9)$$

where P_{decay} is the probability for at least one LDSP to decay inside the radial location of the detector, ϵ_{geo} is the geometric acceptance for the LDSP direction to intersect with the detector, and N_{DS} is the number of DS states produced. For more than one production mode, a sum is implied. Note that we have defined a signal event as one in which at least one LDSP decays inside the detector. The case of more than one LDSPs can be accounted for by multiplying the single LDSP probability with the number of LDSPs produced, and it's included in the definition of P_{decay} (see Appendix C for a detailed discussion of this). As the final step to get the number of signal events S , we have to express N_{DS} in Eq. (9) in terms of the (inclusive) signal cross section σ_S as

$$N_{\text{DS}} = \frac{N_{\text{POT}}}{\sigma_{\text{pN}}} \sigma_S, \quad (10)$$

where σ_S is the cross section for DS production, N_{POT} is the total number of proton delivered on target during the duration of the experiment (projected years for future experiments), and σ_{pN} is the typical proton-nucleus cross section for the proton beam hitting the target, taken constant for the center-of-mass energies of the experiments we consider [60],

$$\sigma_{\text{pN}} = A^{0.77} 49.2 \text{ mb}, \quad (11)$$

with A the target nucleus' atomic weight. In Eq. (10), we are considering only DS production in the first interaction length of the target (or the dump for beam dump experiments), neglecting production happening at later lengths with a degraded beam. Our computations are therefore conservative.

Specific to the case of meson decays, for a given meson \mathbf{M} and in a given decay channel C , N_{DS} is given as

$$N_{\text{DS}} = N_{\text{POT}} N_{\text{M}} \text{Br}_C(\mathbf{M} \rightarrow \text{DS}(+\mathbf{m})), \quad (12)$$

where N_{M} is the number of mesons produced per collision and Br_C is the branching ratio of the meson \mathbf{M} to the DS (which may be in association with other mesons \mathbf{m}).

Strictly speaking, the various factors that go into the estimation of the number of signal events N_{signal} depend on the kinematic information, the production mode, and the details of the detector (e.g., on- vs off-axis). For example,

depending on p_{DS}^2 , the boost of the DS states and therefore its decay probability is different. Further, depending on whether the DS is produced with a nonzero transverse momentum or not, the angle subtended at the detector can be different. The correct procedure would be to consider differential quantities and integrate over the allowed range.² This however can obscure the relation between a given experiment and the probed parameter space. As a way out, we use the average value of boost factor for estimating the probability, and compute the average geometric acceptance. In Appendix E we compare this procedure, referred to as factorized approach, with the exact procedure, called the full approach, and show that the difference between the two is small.

This simplified strategy to compute bounds is useful for the following reason. While production quantities such as the cross section depend in a trivial way on Λ_{UV} and very weakly on Λ_{IR} via the kinematic condition $p_{\text{DS}}^2 \geq n_{\text{LDSP}}^2 \Lambda_{\text{IR}}^2$, the decay probability depends on both parameters. By using averages in the production quantities allows factorizing them from the decay probability. This procedure therefore allows an analytic understanding of Λ_{UV} dependence on the number of signal events. Given the vast array of cases, coming from different experiments, different production channels (which can depend also on extra parameters like the dimension Δ), different decay channels, and the strongly vs weakly coupled scenario, this factorization allows to track the Λ_{UV} dependence clearly, and also speeds up the computations.

We next briefly outline the details of the three production modes discussed earlier.

A. Meson decays

The considered portals between the SM and the DS can cause mesons to decay into DS states. Once the mesons are produced by the incoming proton hitting the target, they can decay in two ways. The first possibility is a heavier meson \mathbf{M} decaying into a lighter SM state (such as another meson \mathbf{m}) along with DS states. This is to be contrasted with the case when the mesons decay just into the DS states and nothing else. These two are the radiative decay and annihilation decay modes respectively. The differential production cross section for radiative decay, as shown in Fig. 1 for the Z portal, is flat in p_{DS} up to kinematic threshold. The decay width for both modes can be approximately estimated, keeping the portal generic,

$$\Gamma \sim \kappa^2 g_{\text{SM}}^2 \Phi(\Delta) \left(\frac{f_{\mathbf{M}}}{M} \right)^a \begin{cases} \frac{v^2 M^{2\Delta} \mathcal{O}^{-1}}{m_h^4 \Lambda_{\text{UV}}^{2\Delta} \mathcal{O}^{-4}}, & \mathcal{O}H^\dagger H \text{ portal,} \\ \frac{M^5}{\Lambda_{\text{UV}}^4}, & JJ \text{ portal,} \end{cases} \quad (13)$$

²Note that for meson annihilation decay $\mathbf{M} \rightarrow \text{DS}$, p_{DS}^2 is fixed to M^2 .




FIG. 2. The underlying quark level transition in DS production via Z portal in flavor-violating decays such as $K \rightarrow \pi + DS$.

where κ is the portal coupling, M is the mass of the parent meson, f_M is the decay constant, $\Delta_{\mathcal{O}}$ is the dimension of \mathcal{O} , g_{SM} is a dimensionless coupling built out of dimensionless SM couplings (like the gauge couplings, loop factors, extra SM particles' phase space, and relevant spurions), and $\Phi(\Delta)$ is the phase space factor coming from the integration over the DS degrees of freedoms [e.g., see Eq. (14)]. The exponent of the dimensionless ratio (f_M/M) depends on the process, and is -2 for processes coming from the axial anomaly, $+2$ for tree-level processes from the chiral Lagrangian and 0 for processes directly proceeding through the portal (without going through the chiral Lagrangian). In this estimate we have ignored the lighter meson mass for radiative decay, and have not included the meson form factors for simplicity. In our full analysis we include all these effect. We next discuss specific details of the radiative and annihilation decays as DS production modes.

I. Radiative decays

For the radiative decay of the form $\mathbf{M} \rightarrow \mathbf{m} + DS$ proceeding via a flavor violating loop, the DS state is produced either by the quark line, and/or by the internal W loop (which is necessary to change the quark flavor). This depends on the portal operator. For $J_{\mu}^{\text{SM}} J_{\text{DS}}^{\mu}$ portal where J_{μ}^{SM} is the quark current, the DS is produced just by the quark lines, whereas for $J_{\mu}^{\text{SM}} = iH^{\dagger} \overleftrightarrow{D}_{\mu} H \sim Z_{\mu}$, the DS states can be produced by attaching a Z to the quarks, or to the W in the loop.³ The DS states can also be produced by the Higgs portal $\mathcal{O}H^{\dagger}H$. To understand their relative importance, let us consider the ratio of the branching ratios of the two different portals for DS production,

$$\frac{\text{BR}_{\text{OHH}}}{\text{BR}_{\text{JHDH}}} \sim \frac{3456\Gamma(\Delta_{\mathcal{O}}+1/2)}{\pi^{1/2}\Gamma(2\Delta_{\mathcal{O}})\Gamma(\Delta_{\mathcal{O}}-1)} \frac{M^2(M-m)^{2\Delta_{\mathcal{O}}-6}}{m_h^4 \Lambda_{\text{UV}}^{2\Delta_{\mathcal{O}}-8}} \quad (14)$$

with M (m) being the mass of the heavy (light) meson, $\Delta_{\mathcal{O}}$ the dimensionality of the operator \mathcal{O} , where for the Higgs portal case we have used the usual effective Lagrangian coupling $C_{ij} \bar{d}_L^i d_R^j h$ between the flavor changing quarks and Higgs (see Ref. [30] for example) and for the Z -portal case, we have considered only the top-quark contribution in the loop of the $\mathcal{O}(m_t^2/m_W^2)$. It is clear that for $\Delta_{\mathcal{O}} \geq 4$, production through the Higgs portal is suppressed with

³We are working in the unitary gauge.

respect to the Z portal and will give weaker bounds. We discuss this in more detail in Sec. V D.

For production through a $\Delta_{\mathcal{O}} = 3$ Higgs portal, even though the Λ_{UV} scale probed is higher than Z portal-production case (discussed in more detail in Sec. V D), the bound is still at most only marginally stronger compared to missing-energy searches at LHC [30]. In this subsection we will mostly focus only on the Z portal production for mesons, but will make some comments about the Higgs portal case in Sec. V.

Examples of radiative meson decay processes are $K^+ \rightarrow \pi^+ + DS$, $B^+ \rightarrow K^+ + DS$, and $D^+ \rightarrow \pi^+ + DS$, and a prototypical diagram (for $K^+ \rightarrow \pi^+ + DS$) is shown in Fig. 2. In general, these processes proceed through insertion of two Cabibbo-Kobayashi-Maskawa (CKM) entries, so that for flavor i going to j , the amplitude approximately scales as $\sum_k V_{ik}^{\text{CKM}} V_{kj}^{\text{CKM}} f(m_k/m_W)$, where m_k is the quark mass of flavor k , m_W is the W mass and $f(x)$ is a loop function [61,62]. For D mesons, for which the underlying process is $c \rightarrow u$, there is no top quark in the loop, as opposed to B, K decays, which makes the D -meson process suppressed. As a result, the D decays are not very constraining—e.g., the large number of D mesons expected at SHiP (enhancement by $\sim 10^4$ compared to B -meson production, see Ref. [63]) is not enough to overcome the CKM suppression of $\sim 10^{-12}$.

Due to the abundant number of K mesons produced at neutrino experiments, $K \rightarrow \pi + DS$ decay is an important mode for DS production. For this process, and for the Z portal case, the Feynman diagrams are shown in Fig. 2. Note that one must include penguin diagrams as well as self-energy diagrams [61,62]. The DS production rates can be obtained from the SM calculation for $d\bar{s} \rightarrow \bar{\nu}\nu$, but with some modifications. We can use the SM results if we keep only the penguin diagrams and omit the box diagrams in the $d\bar{s} \rightarrow \bar{\nu}\nu$ process, since the latter are specific to the neutrino coupling (e.g., see Ref. [62]). This however must be done in the unitary gauge since the box and the penguin diagrams are needed together to make the result gauge invariant in an arbitrary gauge, but their gauge-dependent parts vanish individually in the unitary gauge [62].⁴ Once these subtleties are addressed, we can simply replace the neutrino-current

⁴If we stayed in arbitrary gauge, the DS would also couple to the longitudinal modes of W and hence the box diagrams would also contribute. In the unitary gauge, $H^{\dagger} \overleftrightarrow{D}_{\mu} H \sim Z_{\mu}$, the DS does not couple to W , and the box diagrams' contributions vanish.




FIG. 3. DS produced via Z portal in annihilation decays of vector mesons.

coupling to Z, $g_{EW}/2 \cos \theta_W (\bar{\nu}_L \gamma_\mu \nu_L)$, with the DS current coupling to Z, $(\kappa_J v_{EW} m_Z / \Lambda_{UV}^2) J_\mu^{DS}$. This allows us to write the rate of decay of $K^+ \rightarrow \pi^+ + DS$, using the optical theorem, as

$$\begin{aligned} \Gamma_{K^+ \rightarrow \pi^+ + DS} &= \frac{1}{2M_K} \left(\frac{G_F g_{EW} \cos \theta_W}{\sqrt{2} 8\pi^2} \right)^2 \frac{m_Z^2 v_{EW}^2 \kappa_J^2}{\Lambda_{UV}^4} \\ &\times \left(\sum_{j=c,t} V_{js}^* V_{jd} \bar{D}(x_j, x_u = 0) \right)^2 \\ &\times \int \frac{d^3 p_\pi}{(2\pi)^3} \frac{1}{2E_\pi} \mathcal{M}_\mu \mathcal{M}_\nu^* \\ &\times 2 \text{Im} \langle J_{DS}^\mu(p_{DS}) J_{DS}^\nu(p_{DS}) \rangle, \end{aligned} \quad (15)$$

where G_F is the Fermi constant, θ_W is the weak mixing angle, g_{EW} is the electroweak gauge coupling, M_K is the mass of the K meson, v_{EW} is the Higgs VEV, m_Z is the Z boson mass, $E_\pi = \sqrt{m_\pi^2 + |\vec{p}_\pi|^2}$, $p_{DS} = p_K - p_\pi$, $\mathcal{M}_\mu = \langle \pi^+ | \bar{d} \gamma_\mu s | K^+ \rangle$ is the SM-QCD matrix element (see Appendix B 1 for details), and V_{ij} are CKM matrix elements. The loop functions $\bar{D}(x_j)$, where $x_j = m_j^2/m_W^2$, sum the contributions from various diagrams. The upper limit for p_π integral is fixed by the kinematic requirement $p_{DS}^2 \geq n_{LDS}^2 \Lambda_{IR}^2$.

Apart from the decay $K^+ \rightarrow \pi^+ + DS$, one can also consider the decays of K_L^0 and K_S^0 . We can obtain the partial width of K_L^0 from that of K^+ using Ref. [61] by replacing $|V_{js}^* V_{jd}|^2$ in Eq. (15) with $|\text{Im}(V_{js}^* V_{jd})|^2$. The $K_S^0 \rightarrow \pi^0 + DS$ decay is less constraining since it has a smaller branching ratio due to the large width of K_S^0 (see Refs. [16,64]).

We next consider decays of B mesons to DS which is relevant at proton-beam experiments with higher beam energies (e.g., SHiP and CHARM, with $E_{\text{beam}} \sim 400$ GeV, $\sqrt{s} \sim 27$ GeV). These high-energy proton beam experiments would also have a high K -meson production rate but a large number of them get absorbed in the beam dump or target. Unlike B mesons, kaons have a decay length⁵ which largely exceeds the hadronic-interaction length (l_H); hence they tend to be absorbed in thick targets (for a target length of several l_H s) and only a fraction of them then decay to DS before absorption [55,65]. For estimating this, we use Ref. [65] for SHiP, and Ref. [55] for CHARM.

⁵The decay length of K^\pm, K_L^0 is ~ 3 meters $\gg l_H \sim 15.3$ cm for SHiP and CHARM target [55].

The B -meson decays to lighter mesons like K and π take place via Z -penguin diagrams which we already encountered in the case of $K^+ \rightarrow \pi^+ + DS$ (see Fig. 2) except for the appropriate exchange of external quark flavors ($b \rightarrow s/d + DS$ instead of $s \rightarrow d + DS$). Of all the B decay modes, we find that the largest contribution to signal comes from the decays $B \rightarrow K + DS$ and $B \rightarrow K^* + DS$ [66]. For example, even though the partial width for the decay of $B_s \rightarrow \phi + DS$ is twice of $B \rightarrow K + DS$, the number of signal events from B_s decays are suppressed due to smaller number of B_s mesons produced with respect to B^\pm and B_0 mesons at SHiP [63]. The contribution from $B \rightarrow \pi + DS$ is suppressed with respect to $B \rightarrow K + DS$ by a factor ~ 20 coming from $|V_{ts}|^2/|V_{td}|^2$ that enters in the respective decay widths [61]. This same suppression applies when comparing B meson decays to vector mesons: $B \rightarrow \rho + DS$ is suppressed with respect to $B \rightarrow K^* + DS$. We calculate the partial decay widths for these B decays in the same way as in Eq. (15) using the appropriate QCD matrix elements from Eqs. (B2) and (B4) in Appendixes B 1 and B 2.

We do not consider DS production from radiative decays of pseudoscalar mesons like π, η, η' . Their radiative decay into $\gamma + DS$ through a generic JJ portal is suppressed by the loop factor from the chiral anomaly triangle diagram, from the electromagnetic coupling and from the lightness of the meson in the π case [33].⁶

Radiative decays of vector mesons like ρ and ω can also produce DS via decay modes like $\rho^0 \rightarrow \pi^0 + DS$, etc. These decays would occur via flavor-conserving transitions producing DS either through Z portal or SM vector quark current. The number of DS events from this mode is subleading due to the large width of ρ meson with respect to K meson width ($\Gamma_\rho \sim 10^{-1}$ GeV $\gg \Gamma_K \sim 10^{-17}$ GeV). Moreover, the radiative decays of vector mesons $V \rightarrow DS + P$, where P is a generic pseudoscalar, are suppressed since the interaction mediating the process come from the same triangle diagram mediating pseudoscalar radiative decays like $\pi^0 \rightarrow \gamma + DS$.

Recently [67] considered three-body leptonic decays of mesons to put bounds on leptophilic ALPs. In our case too, DS can be produced from such leptonic charged meson decays such as from the decay $K^+ \rightarrow \mu^+ + \nu + DS$ via Z portal. However, we find this mode to be very suppressed with respect to $K \rightarrow \pi + DS$, due to phase-space suppression (see also [68]).

Eventually, to calculate the number of DS events from a meson decay, we use Eq. (12). It is clear from Eq. (12) that the meson decay mode that gives the strongest bound would depend on N_M , the number of parent mesons produced per

⁶For the Z portal case, one external leg of the triangle diagram would produce Z which can couple with DS. This mode can give bounds at LSND due to the huge number of pions ($N_{\pi_0} \sim 10^{22}$), and we find that the Λ_{UV} probed is comparable to CHARM in the meson-production mode.

POT at a given neutrino experiment. In general, this can be estimated as the ratio of production cross section of the meson to the total cross section between proton beam and target: $N_M = \sigma_{pN \rightarrow M} / \sigma_{pN}$. We take these numbers for various experiments from Ref. [63] (also see references within) for 400 GeV beam energy and from Ref. [14] for 120 GeV beam energy, which are obtained using PYTHIA simulations.

2. Annihilation decays

DS states can also be produced via annihilation decays of vector mesons through the JJ portal; $V \rightarrow DS$ where V can be $\rho, \omega, \phi, J/\psi$ (see Fig. 3). We do not consider DS production from the annihilation decay of pseudoscalar mesons from this portal since it will not be model independent under our approach [30]; the pseudoscalar decay matrix element is proportional to its momentum p_μ , which either vanishes when contracted to a conserved DS current, or gives a term proportional to a new, model dependent scale if the DS current is not conserved (corresponding to the internal DS symmetry breaking scale).

In principle the same topology can happen for the Higgs portal and scalar mesons (the matrix element for the spin-1 annihilation through this portal vanishes). However, given the uncertainties in the details of scalar meson production and their subdominance, we do not consider this possibility here. The leading contribution in this topology for the Higgs portal comes from FCNC CP -violating pseudoscalar annihilation decays such as $K \rightarrow DS$ [69], and we will briefly discuss them together with radiative decays in Sec. V.

For a general V , and for the case of $V \rightarrow DS$ via Z portal, we can compute the decay width as before,

$$\begin{aligned} \Gamma(V \rightarrow DS) &= \frac{1}{2m_V} \frac{1}{3} g_Z^2 \frac{\kappa_J^2 v^2}{m_Z^2 \Lambda_{UV}^4} f_V^2 m_V^2 \\ &\times \sum \epsilon_\mu^*(p) \epsilon_\nu(p) 2 \text{Im} \langle J_{DS}^\mu(p) J_{DS}^\nu(p) \rangle |_{p^2=m_V^2} \\ &= \frac{\kappa_J^2 c_J g_Z^2 v^2 m_V^3 f_V^2}{96\pi m_Z^2 \Lambda_{UV}^4}, \end{aligned} \quad (16)$$

where m_V is the mass of the vector meson, f_V is the decay constant defined by $\langle V(p) | \bar{q} \gamma_\mu q | \Omega \rangle = i f_V m_V \epsilon_\mu^*(p)$, $\epsilon_\mu^*(p)$ is the polarization vector for V meson, and g_Z is the coupling of $\bar{q} \gamma_\mu q$ to the Z -boson. Here we have again used the optical theorem to do integration over DS phase space and used the expressions reported in [30] for the imaginary part of the correlators at $m_V \gg \Lambda_{IR}$. Out of ϕ, ω and J/ψ , the largest branching ratio to DS would be that of J/ψ because of the narrow total width, and a partial width which is enhanced by the mass.

We find that the bounds from $J/\psi \rightarrow DS$ are comparable to those from $B \rightarrow K/K^* + DS$ decays at SHiP. Despite $\text{BR}(J/\psi \rightarrow DS) / \text{BR}(B \rightarrow K/K^* + DS) \sim 10^{-2}$, the large

number of J/ψ mesons expected at SHiP as compared to B mesons, $N_{J/\psi} / N_B \sim \mathcal{O}(100)$ compensates for this.

For neutrino experiments based on the 120 GeV NuMI beam line, annihilation decays of lighter vector mesons like ρ, ω, ϕ can give contribution to signal events. Out of the three vector mesons ρ, ω, ϕ , we find that the leading contribution to DS production is via ϕ meson decay to DS. We can compare the branching ratio for ϕ and ω decay to DS via Z portal,

$$\frac{\text{BR}(\phi \rightarrow DS)}{\text{BR}(\omega \rightarrow DS)} = \frac{\text{BR}(\phi \rightarrow e^+ e^-) m_\phi^4 \left(-\frac{1}{2} + \frac{2}{3} \sin^2 \theta_W\right)^2}{\text{BR}(\omega \rightarrow e^+ e^-) m_\omega^4 e_s^2 4 \sin^4 \theta_W}, \quad (17)$$

where $e_s = -1/3$ is the EM charge of strange quark. Using this we expect DS produced in ϕ decay to dominate over ω decay to DS by a factor given by $N_\phi / N_\omega \times \text{BR}(\phi \rightarrow DS) / \text{BR}(\omega \rightarrow DS) \sim 0.007 / 0.03 \times 50 \sim 10$. Here we have used numbers for ϕ meson production at 80 GeV from Ref. [70] and ω meson production at 120 GeV from [33]. A similar estimate shows that the case of ρ is also subleading. Therefore we only focus on the ϕ decay and do not consider ρ and ω . Note that ρ and ω annihilation decays overlap with the (vector) bremsstrahlung production mode when p_{DS}^2 hits the resonance peak [52], not including them avoids overcounting such contributions. We do not consider the annihilation decays of heavier mesons like Υ since its production will be very suppressed at neutrino experiments due to its large mass.

Now we outline how we compute the LDSP boost entering the decay probability and geometric acceptance factors for the meson production mode. More details can be found in Appendix D 1. In order to calculate the decay probability of the LDSP, we use the following estimate for the average boost factor for the LDSP produced from meson decays:

$$\langle \gamma \rangle_{\text{LDSP}} \approx \frac{\langle E_{DS}^{\text{lab}} \rangle}{\langle n_{\text{LDSP}} \rangle \Lambda_{IR}}, \quad (18)$$

where $\langle E_{DS}^{\text{lab}} \rangle$ is the average energy of the DS produced from parent meson decay in the lab frame. We have checked that an honest average of $\langle \gamma \rangle_{\text{LDSP}}$ matches this estimate very well. To obtain $\langle E_{DS}^{\text{lab}} \rangle$, the strategy is as follows: For radiative decays of the form $M \rightarrow m + DS$, in the parent meson rest frame, the DS 3-momentum $\vec{p}_{DS}^0 = \{(p_{DS}^0)_T, (p_{DS}^0)_z\}$ can be written using energy conservation as

$$|\vec{p}_{DS}^0| = \sqrt{\frac{(M^2 - m^2 + p_{DS}^2)^2}{4M^2} - p_{DS}^2}, \quad (19)$$

and fixes $(p_{DS}^0)_z = |\vec{p}_{DS}^0| \cos \theta_{DS}^0$, where θ_{DS}^0 is the angle that the DS makes with the meson flight direction, in its rest frame. For annihilation decays, of the form $M \rightarrow DS$, the DS

TABLE I. Average quantities $\langle\gamma\rangle_{\text{DS}} = \langle E_{\text{DS}}^{\text{lab}}/p_{\text{DS}}\rangle$, $\langle\gamma\rangle_{\text{LDSP}}$, and $\langle\epsilon_{\text{geo}}\rangle$ for Z portal and H portal production, for various production modes, and for the weak/strong case. The reported numbers are for fixed $\Lambda_{\text{IR}} = 10$ MeV. For Higgs portal, we have taken $\Delta_{\mathcal{O}} = 4$. The shown numbers are for DUNE-MPD (at 120 GeV) and SHiP (at 400 GeV) target materials (which sets the target atomic weight and number A , Z respectively). The average DS boost $\langle\gamma\rangle_{\text{DS}}$ depends on the weak/strong case through the kinematic condition $\sqrt{p_{\text{DS}}^2} \geq n_{\text{LDSP}}\Lambda_{\text{IR}}$ imposed when calculating the average, and is a weak dependence. Annihilation decays of vector mesons does not proceed through the Higgs portal due to mismatch in quantum numbers.

$E_{\text{beam}}^{\text{lab}}$ (GeV)	Z-portal production						H-portal production					
	$\langle\gamma\rangle_{\text{DS}}^{\text{weak}}$	$\langle\gamma\rangle_{\text{DS}}^{\text{strong}}$	$\langle\gamma\rangle_{\text{LDSP}}^{\text{weak}}$	$\langle\gamma\rangle_{\text{LDSP}}^{\text{strong}}$	$\langle\epsilon_{\text{geo}}\rangle^{\text{weak}}$	$\langle\epsilon_{\text{geo}}\rangle^{\text{strong}}$	$\langle\gamma\rangle_{\text{DS}}^{\text{weak}}$	$\langle\gamma\rangle_{\text{DS}}^{\text{strong}}$	$\langle\gamma\rangle_{\text{LDSP}}^{\text{weak}}$	$\langle\gamma\rangle_{\text{LDSP}}^{\text{strong}}$	$\langle\epsilon_{\text{geo}}\rangle^{\text{weak}}$	$\langle\epsilon_{\text{geo}}\rangle^{\text{weak}}$
Drell-Yan												
120 (DUNE-MPD)	12	12	1490	160	0.004	0.004	9	9	1600	150	0.002	0.002
400 (SHiP)	25	25	4310	400	0.63	0.63	17	17	4830	364	0.54	0.54
Dark bremsstrahlung												
120 (DUNE-MPD)	80	80	4500	655	0.040	0.040	74	74	4630	650	0.039	0.039
400 (SHiP)	270	270	15000	2200	1	1	250	250	15700	2200	0.96	0.97
Meson radiative decay $K \rightarrow \pi + \text{DS}$												
120 (DUNE-MPD)	31	26	215	48	0.003	0.004	31	26	215	48	0.003	0.004
400 (SHiP)	55	45	375	84	0.79	0.89	55	45	375	84	0.79	0.89
Meson annihilation decay $\phi \rightarrow \text{DS}$												
120 (DUNE-MPD)	8	8	403	61	0.001	0.001
400 (SHiP)	14	14	702	107	0.27	0.27

3-momentum in the meson rest frame is zero by momentum conservation. We further assume that 3-momentum of the mesons that decay to DS is perfectly aligned along the beam axis i.e., $\theta_{\text{meson}} = 0$.⁷

We next calculate $[E_{\text{DS}}^{\text{lab}}, (p_{\text{DS}}^{\text{lab}})_z]$ from $[E_{\text{DS}}^0, (p_{\text{DS}}^0)_z]$ using the boost and the velocity of the parent meson in the lab (which is along the z -axis), obtained from the average meson momentum values for various experiments from Table 6 in Ref. [33]. From $(p_{\text{DS}}^{\text{lab}})_z$, we can obtain $|\vec{p}_{\text{DS}}^{\text{lab}}|$ by noting that the transverse component is unaffected by the z -direction boost, so that everything is a function of θ_{DS}^0 and p_{DS}^2 . Finally, to get the average value of DS 3-momentum $\langle|\vec{p}_{\text{DS}}^{\text{lab}}|\rangle$, we average over $\cos\theta_{\text{DS}}^0$, since DS is isotropic in this variable and set p_{DS}^2 to its average value for each radiative meson case. We have again checked that this matches a true average.

To get the final number of signal events as in Eq. (9), we also need the geometric acceptance, which we again compute as an average. See Appendix D 1 for details of these computations, and Appendix E for a comparison between using this average procedure with a more refined analysis. Some typical values of $\langle\gamma\rangle_{\text{LDSP}}$ are given in Table I.

⁷A more refined analysis using [14] shows that the most probable value for the ratio between transverse and longitudinal components of 3-momentum of decaying K mesons $|p_{\text{meson}}^{\text{T}}/p_{\text{meson}}^z| \sim \theta_{\text{meson}} \sim 10^{-2} \ll 1$. Using a nonzero but small value of θ_{meson} does not change our final results.

B. Drell-Yan production

If the typical exchanged momentum from the protons to the DS is comparable or larger than Λ_{QCD} , the process is able to probe the partonic constituents of the nucleon. Given the energy scales involved, the protons are ultra-relativistic, and using the PDF language to model the interaction between the constituents is justified. Notice that in our case, the condition to probe the partonic structure of the nucleon is $p_{\text{DS}}^2 \gtrsim 1 \text{ GeV}^2$, which is a request on the total DS system, and not on the mass Λ_{IR} of the DS constituents. This is unlike what happen in models in which the mediator is produced on shell, such as in light dark photon models. The production cross section is in general dependent on the portal. A general estimate for the amplitude of DY through a given portal can be obtained on dimensional grounds, by assuming the typical momentum to be $\sqrt{p_{\text{DS}}^2}$, and integrating over it to get the cross section. For the Higgs portal, the partonic cross section comes from Higgs exchange and is dominated by gluon initial states, while for Z portal, there is a Z exchange, and the initial states are the quarks. For the Z-aligned JJ portal, the results of Z portal apply, once appropriately rescaled, if the couplings are assumed to be Z-aligned (both in axial-vector and isospin space). The Feynman diagram for such a process is shown in Fig. 4. Due to the similarity with DY annihilation process we dub this production channel DY.

Consider first the Higgs portal. The leading interaction at the constituent level is due to gluon-gluon fusion (ggF) processes: indeed light quarks, while abundant in the proton, have a suppressed coupling to the Higgs, while




FIG. 4. Prototypical Drell-Yan process in Z portal for the DS production.

heavy quarks are rare in the proton. Following [66], the effective ggF operator is, after integrating out the Higgs,

$$\mathcal{L} \supseteq F(\hat{s}) \frac{\kappa_{\mathcal{O}}}{\Lambda_{\text{UV}}^{\Delta_{\mathcal{O}}-2}} \frac{\alpha_s}{16\pi m_h^2} G^{\mu\nu a} G_{\mu\nu}^a \mathcal{O}, \quad (20)$$

where $F(\hat{s})$ is a function of the center-of-mass energy \hat{s} of the process that accounts for the loops of internal quarks. Given that our computation is valid only for p_{DS} above the QCD scale, we retain in F only the contributions coming from the top, bottom, charm, and strange quarks. This expression holds for center-of-mass energies much smaller than the Higgs mass (true for typical neutrino and beam dump experiments).

The cross section to produce a DS shell of total momentum p_{DS}^2 can be computed by integrating over the DS phase space using optical theorem,

$$\begin{aligned} \sigma^{\text{DY(Higgs)}} &= A \sigma_{\text{pp}}^{\text{DY(Higgs)}} = \frac{A \alpha_s^2}{1024 m_h^4 \pi^{7/2}} \frac{\kappa_{\mathcal{O}}^2 c_{\mathcal{O}}}{\Lambda_{\text{UV}}^{2\Delta_{\mathcal{O}}-4}} \frac{\Gamma(\Delta_{\mathcal{O}}+1/2)}{\Gamma(2\Delta_{\mathcal{O}})\Gamma(\Delta_{\mathcal{O}}-1)} \\ &\times \int_{Q_0^2}^s d p_{\text{DS}}^2 |F(p_{\text{DS}}^2)|^2 \frac{p_{\text{DS}}^{2\Delta_{\mathcal{O}}-4}}{\Lambda_{\text{UV}}^{2\Delta_{\mathcal{O}}-4}} \\ &\times \int_{p_{\text{DS}}^2/s}^1 \frac{dx}{sx} f_g(p_{\text{DS}}, x) f_g(p_{\text{DS}}, p_{\text{DS}}^2/(sx)), \quad (21) \end{aligned}$$

where A is the atomic number of the target nucleus, f_g are the gluon PDFs,⁸ x is the longitudinal momentum fraction of one of the initial gluons in the CM frame and \sqrt{s} is the center-of-mass energy of the protons. The lower limit of the integral over p_{DS}^2 is cutoff at $Q_0^2 = (1.3 \text{ GeV})^2$, the lowest for which the PDFs have been fitted, and below which the process does not probe a single parton and the DY picture breaks down. For consistency, Q_0^2 must be more than the minimum invariant mass of the DS system, $(n_{\text{LDSP}} \Lambda_{\text{IR}})^2$, which we impose internally.

⁸To compute the PDF integral, we used the nCTEQ15 PDF values [71], included in the MANEPARSE *Mathematica* package [72].

Next consider the Z portal. Since the quark-Z coupling depends only on the up or down type of the quark, the process is dominated by light quark-antiquark annihilations. We will consider only contributions coming from up and down quarks, for which the couplings are given as

$$\begin{aligned} g_u^2 &= \frac{g_{\text{EW}}^2}{\cos^2 \theta_W} \left(\frac{1}{8} + \frac{4}{9} \sin^4 \theta_W - \frac{1}{3} \sin^2 \theta_W \right), \\ g_d^2 &= \frac{g_{\text{EW}}^2}{\cos^2 \theta_W} \left(\frac{1}{8} + \frac{1}{9} \sin^4 \theta_W - \frac{1}{6} \sin^2 \theta_W \right). \quad (22) \end{aligned}$$

Unlike the ggF case, the relevant PDFs depend on whether the target nucleon is a proton or a neutron. We approximate the neutron PDFs f_i^n to be the isospin-rotated PDFs of the proton f_i^p ,

$$f_u^p = f_d^n, \quad f_d^p = f_u^n, \quad f_u^n = f_d^p, \quad f_d^n = f_u^p. \quad (23)$$

The partonic cross section for the pp/pn interaction in the limit of massless quarks reads,

$$\begin{aligned} \sigma_{pp/pn}^{\text{DY(Z)}} &= \frac{1}{1152\pi} \frac{g_{\text{EW}}^2 v^4 \kappa_J^2 c_J}{m_Z^4 \cos^2 \theta_W \Lambda_{\text{UV}}^4} \int_{Q_0^2}^s d p_{\text{DS}}^2 p_{\text{DS}}^2 \\ &\times \int_{p_{\text{DS}}^2/s}^1 \frac{dx}{sx} X_{pp/pn}(s, x, p_{\text{DS}}^2), \quad (24) \end{aligned}$$

where

$$\begin{aligned} X_{pp}(s, x, p_{\text{DS}}^2) &= 2 \sum_{i=u,d} (g_i^2 f_i^p(x) f_i^p(p_{\text{DS}}^2/(sx))), \\ X_{pn}(s, x, p_{\text{DS}}^2) &= g_u^2 f_u^p(x) f_d^n(p_{\text{DS}}^2/(sx)) \\ &\quad + g_u^2 f_u^n(x) f_d^p(p_{\text{DS}}^2/(sx)) \\ &\quad + g_d^2 f_d^p(x) f_u^n(p_{\text{DS}}^2/(sx)) \\ &\quad + g_d^2 f_d^n(x) f_u^p(p_{\text{DS}}^2/(sx)). \quad (25) \end{aligned}$$

In the PDFs used, we have taken the factorization scale to be the exchanged momentum p_{DS}^2 and not indicated it explicitly to keep the expressions simpler. Putting the contributions from the protons and neutrons together, the total DY cross section for the Z portal is

$$\sigma^{\text{DY(Z)}} = Z \sigma_{pp}^{\text{DY(Z)}} + (A - Z) \sigma_{pn}^{\text{DY(Z)}}, \quad (26)$$

where Z, A are respectively the atomic and weight number of the target nuclei. Notice that in both Higgs- and Z-portal scenarios the cross section increase with p_{DS} , as expected on dimensional grounds. The drop at high p_{DS} seen in Fig. 1 is due to the PDF convolutions.

In order to estimate the decay probability, we estimate the average boost of the LDSP in the lab frame (not to be confused with the boost of the total DS system) as given in




FIG. 5. DS produced via Z portal in the bremsstrahlung production mode.

Eq. (18). The value for these averaged quantities is given in Table I.

In principle, the boost should take into account the angle in the DS frame: while particles in the DS frame have roughly the same energy, in the lab-frame particles emitted along the beam are more boosted with respect to particles emitted in the opposite direction. We have checked that this effect is negligible, when restricting to particles hitting the detector. For the geometric acceptance, we notice that in the DY production mode, the DS system has no transverse momentum and is collinear to the beam axis. After boosting the LDSP momentum in the DS frame we compute the angles that corresponds to the detector. To estimate $\epsilon_{\text{geo}}^{\text{DY}}$, we follow the prescription given in Appendix D.

In general, in the DY production mode, events are produced with larger p_{DS}^2 compared to other modes (see Fig. 1). We also find that the average energy of the DS system in the lab frame is not as high as in bremsstrahlung. These lead to lower γ of the DS system, a larger LDSP angular spread and therefore a slightly smaller ϵ_{geo} for on-axis detectors.

C. Dark bremsstrahlung ($pp \rightarrow \text{DS} + X$)

Another possibility is for the DS states to be produced directly from proton as an initial state radiation. In this case, the exchanged momentum p_{DS}^2 is not hard enough to probe the partonic structure. Following [73], we model the process using the initial-state radiation (ISR) splitting function formalism. The idea is to treat the DS as incoming from the leg of the initial beam proton, which then becomes slightly virtual—an almost on shell particle participating in the rest of the process. In the following, we will use the standard jargon; p_T for the transverse momentum of the DS system (in the plane orthogonal to the beam direction) and zE_{beam} for its longitudinal momentum, where E_{beam} is the beam longitudinal momentum (in the lab frame).

The splitting function formalism works well when the virtual particle is almost on shell. This means that in order to get reasonable cross sections, we must integrate the variables p_T and z in a subregion of their kinematically allowed values, in which the virtual proton is not too off shell.

Denoting the proton after DS emission as p' , concretely, we will consider the region in which the virtuality is small,

$$\frac{p'^2 - m_p^2}{E_{p'}^2} = \frac{z^2 m_p^2 + (1-z)p_{DS}^2 + p_T^2}{z(1-z)^2 E_{\text{beam}}^2} < 0.1. \quad (27)$$

The choice of 0.1 is arbitrary and our results are not sensitive to small changes in this. In order to compute the splitting functions, we need to compute the vertex between the proton and the Z or Higgs. For the Higgs, by using low-energy theorems [74] we can compute the coupling between the Higgs and nucleon at zero momentum to be $g_{hNN} h \bar{N} N$, where $g_{hNN} = 1.2 \times 10^{-3}$. To model the momentum dependence of the form factor, we employ a generalization of the extended vector-meson dominance (eVMD) model, in which the DS state interacts with the hadron by mixing with the scalar, CP even hadronic resonances. The resonances' propagators are taken to be Breit-Wigners (BW), and the mixing coefficients are fixed by using sum rules and by fitting the zero momentum values. The form factor is taken from [73] and the specific values used are reported in Appendix A 1.

We also need to take into account the fact that for too high virtuality the quasireal proton stops interacting with the target proton as a coherent object, and the bremsstrahlung computation breaks down. To do this we multiply the previous form factor by a smooth cutoff [75],

$$F_D(Q^2) = \frac{\Lambda_p^4}{\Lambda_p^4 + Q^4}, \quad (28)$$

where $\Lambda_p = 1.5$ GeV is the cutoff, taken to be near the proton mass, and Q^2 is the virtuality of the intermediate proton,

$$Q^2 = \frac{z^2 m_p^2 + (1-z)p_{DS}^2 + p_T^2}{z^2}. \quad (29)$$

Finally, the cross section for the process is calculated by factorizing the total cross section into the bremsstrahlung part and a proton-nucleus (after bremsstrahlung) part. The proton-nucleus cross section $\sigma_{pN}^{\text{nTSD}}$ is calculated using the difference between the total inelastic proton-nucleus cross section and the target single diffractive (TSD) contribution, in which the target nucleus is diffracted but not disintegrated. This choice allows neglecting possible interference between the initial-state and the final-state radiation [60,73]. According to [60,76], $\sigma_{pN}^{\text{nTSD}}$ is a slowly varying function of energy, and for the energies involved, we can approximate it as a constant

$$\sigma_{pN}^{\text{nTSD}} = 762(A/56)^{0.71} (1 - 0.021(56/A)^{0.36}), \quad (30)$$

where A is the target atomic weight. We will now generalize the results of [66,73] to higher-dimensional portals.

So far the discussion applies to any of the portals. However, once a portal is specified, the involved form factors change. Consider first the Higgs portal. Putting everything together, the inclusive production cross section is given as

$$\begin{aligned} \sigma_{pN}^{\text{Brem(Higgs)}} &= \sigma_{pN}^{\text{nTSD}} g_{hNN}^2 \frac{v^2}{16\pi^{9/2} m_h^4} \frac{\Gamma(\Delta_O + 1/2)}{\Gamma(\Delta_O - 1)\Gamma(2\Delta_O)} \\ &\times \frac{\kappa_O^2 c_O}{\Lambda_{\text{UV}}^{2\Delta_O - 4}} \int dp_{\text{DS}}^2 dp_T^2 dz |F_H|^2 \\ &\times z((2-z)^2 m_p^2 + p_T^2) \\ &\times \left(\frac{1}{m_p^2 z^2 + (1-z)p_{\text{DS}}^2 + p_T^2} \right)^2 p_{\text{DS}}^{2\Delta_O - 4}, \quad (31) \end{aligned}$$

where F_H is the Higgs bremsstrahlung form factor built as outlined before, and can be found in Appendix A 1. The limits of integration are chosen to respect the kinematic condition $p_{\text{DS}}^2 \geq m_{\text{LDSP}}^2 \Lambda_{\text{IR}}^2$.

Next consider the Z portal case. The only difference is in the form factors of the axial and vector current of the proton. For the vector case, the production cross section is given as

$$\begin{aligned} \sigma_{pN}^{\text{Brem(Z, Vector)}} &= \sigma_{pN}^{\text{nTSD}} \frac{\kappa_J^2 c_J}{2^{11} \pi^4} \frac{v^4}{m_Z^4 \Lambda_{\text{UV}}^4} \frac{g_{\text{EW}}^4}{\cos^4 \theta_W^4} \\ &\times \int dp_{\text{DS}}^2 dp_T^2 dz |F_Z^V|^2 \\ &\times \left(\frac{2}{z} + \frac{4p_{\text{DS}}^2 z (p_T^2 + m_p^2 (z^2 + 2z - 2))}{(m_p^2 z^2 + (1-z)p_{\text{DS}}^2 + p_T^2)^2} \right). \quad (32) \end{aligned}$$

The vector form factor $F_Z^V(p_{\text{DS}}^2)$ is modeled by ρ (isotriplet) and ω (isosinglet) exchange. We take three states for each tower. Details are given in Appendix A. For the axial case, the cross section is given as

$$\begin{aligned} \sigma_{pN}^{\text{Brem(Z, Axial)}} &= \sigma_{pN}^{\text{nTSD}} \frac{\kappa_J^2 c_J}{2^{11} \pi^4} \frac{v^4}{m_Z^4 \Lambda_{\text{UV}}^4} \frac{g_{\text{EW}}^4}{\cos^4 \theta_W^4} \\ &\times \int dp_{\text{DS}}^2 dp_T^2 dz |F_Z^A|^2 \\ &\times \frac{2}{z} \left(\frac{1}{m_p^2 z^2 + (1-z)p_{\text{DS}}^2 + p_T^2} \right)^2 \\ &\times (p_{\text{DS}}^2 (1-z)^2 + (p_T^2 + z^2 m_p^2)^2 + 2p_{\text{DS}}^2 \\ &\times (m_p^2 z^2 (5 + (z-5)z + p_T^2 (1 + z^2 - z))))). \quad (33) \end{aligned}$$

Similar to the vector case, for the axial form factor we take the respective isotriplet axial vector exchange (there is no contribution from the axial isosinglet resonances). Details about the axial form factor $F_Z^A(p_{\text{DS}}^2)$ are in Appendix A. Combining the vector and the axial pieces we get

$$\sigma_{pN}^{\text{Brem, (Z)}} = \sigma_{pN}^{\text{Brem, (Z, Vector)}} + \sigma_{pN}^{\text{Brem, (Z, Axial)}}. \quad (34)$$

Notice that in Eq. (34) the interference term between the vector and axial piece vanishes, due to the different quantum numbers under parity of the two possible states.

The vector contribution to the cross section is subdominant with respect to the axial one, due to an accidental cancellation in the vectorial quark coupling. The vector contribution has a more narrow distribution in p_{DS}^2 than the axial one, and it is peaked at m_ω^2 ; this is because ω is much more narrow than the isotriplet vectors and axial vectors resonances mixing with the Z . Notice that the same exact computation holds for a JJ portal aligned (in both Lorentz and flavor space) to the Z quantum numbers. For different coupling structure of J_{SM} for a generic JJ portal, we can decompose the proton-vector and axial form factors in their isosinglet and isotriplet components to get the correct form factor, shown in Appendix A.

An estimate for the cross section can be given by exploiting the fact that the cross section is dominated by the Breit-Wigner (BW) peaks in the form factors. The p_{DS}^2 integral of the BW associated with an intermediate meson \mathbf{m} can be estimated as $\pi f_{\mathbf{m}}^2 m_{\mathbf{m}}^3 / (2\Gamma_{\mathbf{m}})$, where we have used the notation of Appendix A. Therefore, using the splitting function to give the correct momentum scaling, up to an $\mathcal{O}(1)$ factor, the cross section for radiating DS particles through a dimension D portal made of a SM operator and a dimension Δ_O DS operator can be estimated as

$$\sigma \sim \sigma_{pN} g_{\text{SM}}^2 \Phi(\Delta) f_{\mathbf{m}}^2 \begin{cases} \frac{v^2}{m_h^4} \frac{m_{\mathbf{m}}^{2\Delta_O - 1}}{\Lambda_{\text{UV}}^{2\Delta_O - 4} \Gamma_{\mathbf{m}}} & : \mathcal{O}H^\dagger H \text{ portal,} \\ \frac{m_{\mathbf{m}}^5}{\Lambda_{\text{UV}}^4 \Gamma_{\mathbf{m}}} & : JJ \text{ portal,} \end{cases} \quad (35)$$

where, as in the meson case, g_{SM} is a dimensionless SM factor built out of dimensionless couplings (like the gauge couplings), $\Phi(\Delta)$ is the phase space factor coming from the integration over the DS degrees of freedoms, $f_{\mathbf{m}}$ is the coupling of proton to the meson \mathbf{m} and $\Gamma_{\mathbf{m}}, m_{\mathbf{m}}$, is its decay width and mass respectively (see Appendix A). In presence of multiple resonances, the estimate can be done by restricting to the leading contribution of the BW. For both the Higgs and Z portals the DS is produced as a collimated state forming an angle with the beam $\theta_{\text{DS}} \sim p_T / E_{\text{DS}}^{\text{lab}} = p_T / (zE_{\text{beam}})$. The average acceptance $\epsilon_{\text{geo}}^{\text{brem}}$ is computed by averaging over all the kinematic variables, and it does not differ much from the one obtained by replacing θ_{DS} with its average. Details of the computation of the geometric acceptance are given in

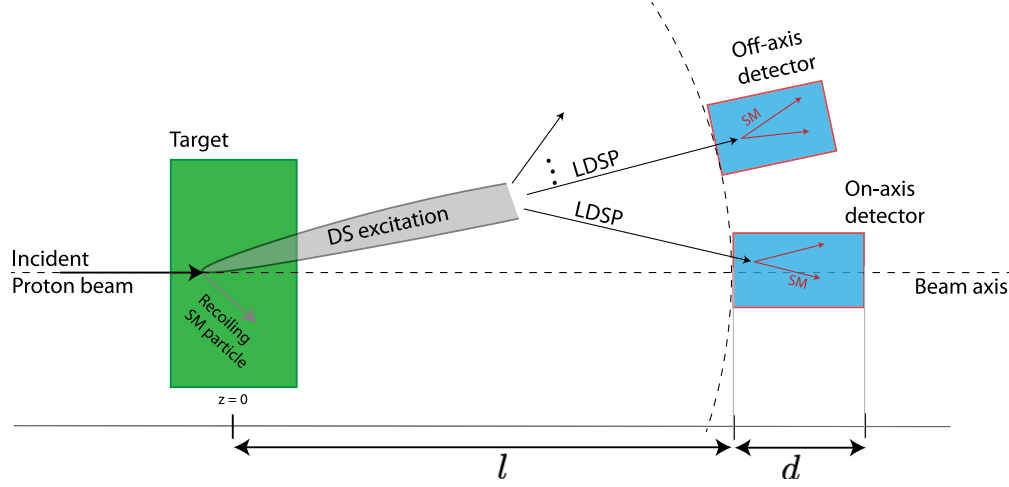


FIG. 6. A cartoon of DS events produced at a typical neutrino detector (drawn not to scale); both on-axis and off-axis cases are shown. In general the DS state may not be produced along the beam axis, as shown, though this is dependent on the production mode; annihilation decays of mesons and DY production modes produce DS along the beam axis, while radiative decays of mesons and DB production mode produce DS at a small angle from the beam direction.

Appendix D. To estimate the average decay probability, we use the average LDSP boost, as defined in Eq. (18) (see also Table I for typical values), with the only difference being that the probability distribution is given by the splitting function. The typical energy of the DS system in the lab frame is roughly $(3/4)E_{\text{beam}}$ for all the portals considered, larger than in DY case.

IV. EXPERIMENTAL SETUPS, SIGNAL, AND BACKGROUND ESTIMATION

In this section we briefly discuss the experiments we consider for obtaining the bounds on the Λ_{UV} and Λ_{IR} scales, and the assumptions we make when obtaining these bounds.

As explained in Sec. II, the relevant characteristics of high intensity experiments are their beam energy E_{beam} , their integrated luminosity (reported as the total protons on target N_{POT}) and geometric details of the experimental setup. Further, the detectors in these experiments can be

placed on-axis (i.e., along the line of the incoming beam) or off-axis (see Fig. 6 for a cartoon of the experimental setup), which can change the geometric acceptance if certain production modes are forward peaked. Specific to the MD production mode, the number of meson N_M produced at a given experiment is an additional input, as seen from Eq. (12). This depends on the target details as well as the energy in the center-of-mass frame, $\sqrt{s} \approx \sqrt{2E_{\text{beam}}m_p}$.

In this work we consider a representative set of past, present, and future experiments. Table II gives the list of considered experiments with the relevant parameters.

We show bounds from recasts of BSM search results from past and current experiments, and projections from current and upcoming experiments, considering a few representatives from each category. We will also consider a future dedicated LLP search experiment, SHiP, for comparison, since it has a broad reach in typical DS models (e.g., see Ref. [83]). We will take the SHiP

TABLE II. The relevant parameters for the experiments considered in this work. The quantities l and d are defined in Fig. 6. θ_{det} stands for the position of the detector center with respect to the beam line, with the origin taken at the interaction point. Entries with zero θ_{det} indicate that the detector is placed along the beam axis. θ_{acc} stands for the detector half-angular opening. Note that for MicroBoone KDAR analysis, the K mesons are produced at rest (in the lab frame) so that θ_{det} is irrelevant. The angle θ_{acc} for this case is measured with the origin at the NuMI hadron absorber, placed $\mathcal{O}(600)$ m from the interaction point [77].

Experiment	N_{POT} (total)	E_{beam} (GeV)	l (m)	d (m)	Off-axis angle, θ_{det} (rad)	θ_{acc} (rad)
CHARM [78–80]	2.4×10^{18}	400	480	35	0.01	0.003
NO ν A-ND [47,81]	3×10^{20}	120	990	14.3	0.015	0.002
MicroBooNE (KDAR) [77]	1.93×10^{20}	120	100	10.4	...	0.013
ICARUS-NuMI [11,16]	3×10^{21}	120	803	19.6	0.097	0.005
DUNE-MPD [14,82]	1.47×10^{22}	120	579	5	0	0.004
SHiP [65,83]	2×10^{20}	400	64	50	0	0.078

parameters to be very optimistic, to have a conservative comparison with the neutrino experiments. A more comprehensive analysis that also considers other dedicated LLP experiments can be useful, and will be done in future. The best bounds from the high intensity experiments that we consider are from the currently under construction DUNE experiment and the proposed SHiP experiment. We also show bounds coming from other experiments running at the same beam energies of these two (120 GeV and 400 GeV, respectively).

Out of past experiments, we show CHARM, a beam dump experiment that ran on the CERN SPS (400 GeV) beamline in the 1980s. CHARM searched for the decay of axionlike particles (ALPs) into a pair of photons, electrons and muons and found no events [78], and we will recast this search for our bounds. Our choice of the past experiments are representative, and not based on the strongest bounds, but rather on considering similar beam energies as DUNE and SHiP (e.g., PS191 and ν CAL, with beam energy of 19.2 GeV and 70 GeV respectively, can give a slightly better bound than CHARM, but are much weaker than the advocated future experiment DUNE. We will comment in Sec. V about these two experiments).

We do not recast heavy neutral lepton decay searches of CHARM or other past experiments like BEBC, or dark particle scatterings as done for example in [84,85], given the different final-state topology.

From existing experiments, we choose the MicroBooNE and ICARUS experiments, based on the 120 GeV NuMI beamline. These are two of the three detectors of Fermilab's short-baseline neutrino program (SBNP) [11].⁹ For MicroBooNE, we will use the analysis in Ref. [77] for dark scalars decaying into electron-positron pairs. For ICARUS, we will use the results in [16], which studied DS coupled through the renormalizable Higgs portal.¹⁰ DM searches at another current experiment, MiniBooNE, based on 8 GeV BNB beamline, use scattering [47,86], and as explained in Sec. II, they require additional model dependent assumptions, so we will not consider them here. Another currently running NuMI-based experiment is NO ν A. We are not aware of any search for DM decays done at this experiment (for a scattering analysis, see Refs. [47,87], based on [81]). Since NO ν A is currently running, we show a possible prospect of such a search. We assume that it will be possible

⁹The third detector, SBND, is too far off-axis with respect to the NuMI beamline and therefore its geometric acceptance is too low to give meaningful constraints. These three detectors also run on the 8 GeV booster neutrino beam (BNB), which is at a lower energy than the NuMI energy, 120 GeV. We find the bounds to be subleading compared to DUNE (but better than other detectors), and do not consider it.

¹⁰Since the target specifications for NuMI beamline experiments are the same as that of proposed DUNE-LBNF beamline, we recycle the meson production numbers for DUNE-LBNF [14] also for the NuMI beamline experiment ICARUS.

to reduce the backgrounds to negligible amounts, given the good angular resolution of the detector.

For future experiments we look at the Deep Underground Neutrino Experiment (DUNE). Reference [14] has proposed the use of the multipurpose, high pressure gaseous chamber—the multipurpose detector (MPD) present in DUNE near detector complex for DS searches. We show projections for our DS scenario for the future DUNE-MPD as well.

Another class of future and existing experiments that are worth considering are the ones that are built in a region directly forward of the beam interaction point at LHC, like the already running FASER [88], FASER ν [89] (see Ref. [90] for the recent analysis) and SND [91], and the proposed experiments at the Forward Physics Facility (FPF) [92]. These experiments are in spirit similar to proton beam dumps, given that unlike traditional collider experiments they are put directly in the forward region, but have a much higher-beam energy. This can potentially be relevant for nonrenormalizable portals. The reach of these detectors have been thoroughly studied in the recent past in the context of dark sectors; a nonexhaustive list of works includes [64,93–95] for renormalizable portals at FASER, [96] for ALPS at FASER, [97] for SND prospects, [98] for a strongly coupled dark sector at the FPF, [99,100] for renormalizable portals at the FPF and [66] for FASER2 prospects. We will briefly comment on their reach on these models in Sec. V.

In order to estimate the sensitivity of the selected current and future experiments, an assessment of the background is needed. We assume that beam dump experiments can be made background free by imposing cuts with $\mathcal{O}(1)$ signal efficiencies, as seen in past searches, e.g., at CHARM [78]. On the other hand, at neutrino experiments, the neutrino beam itself can be a source of background events. At these experiments the typical mass of the LDSPs probed is $\mathcal{O}(10\text{--}100)$ MeV, therefore the available channels for the LDSP decay to the SM are mostly photons and electron-positron pairs, which produce electromagnetic showers in the detector. Heavier decay products, such as muons, will be reconstructed as tracks (but for all practical purposes, we will treat them similar to the showers in this section). The following discussions hold for any of the decay products.

In principle, the two shower signature has no irreducible background. Reducible background events come from hard radiation of a single photon, or from neutral-current ν scattering against a nucleus producing a π^0 , which then decays into $\gamma\gamma$. The produced photons then can convert into e^+e^- pairs, that mimic the signal. However, it's not guaranteed that the two daughter particles will be reconstructed as separate showers.

The typical condition in order to reconstruct the two particles involves an isolation cuts between the decay products, or in other words an angular separation cut.

The specific implementation depends on the specific detector and analysis strategy. We will briefly review what has been suggested in previous works. However, many of the relevant aspects can be understood more generally, which we will elaborate with a relevant prototypical experiment in mind.

For the ICARUS experiment, as suggested in [16], an angular separation of 10° is enough to be able to separate the two showers. Background events instead have a narrow separation between the charged particles, or potentially two showers that do not originate from the same vertex. The angular cut reduces the background events to a negligible amount. In [46] the authors elaborate on an analysis with less stringent cuts but with $\mathcal{O}(100)$ background events. Indeed a strong isolation cut has low efficiency for lighter, and therefore more boosted, LDSPs. This is especially true for the models under consideration here, in which DB and DY production modes are non-negligible and generate LDSPs more boosted than the ones coming from meson decays. For example, in DB, for the weakly coupled case $n_{\text{LDSP}} = 2$, $\Lambda_{\text{IR}} = 100$ MeV, we expect the daughter particles to be separated by an angle of $1/\gamma_{\text{LDSP}} \approx 0.15^\circ$, which is smaller or comparable to the angular resolutions of some of the detectors. This highlights a potential problem in our framework, when reconstructing the signal. For this reason, we suggest that at ICARUS, it might be better to avoid a stringent cut in angular separation and work with $\mathcal{O}(100)$ background events [46], possibly reduced with an energy cut and a cut on the direction of the DS system with respect to the beam. Interestingly, for strongly coupled DS we expect the angular separation condition to be less stringent on the signal. Because of a larger n_{LDSP} , the energy is split among more LDSPs, leading to a suppression of the single LDSP boost factor. For such sectors, assuming an average n_{LDSP} of $\mathcal{O}(10)$, the average separation angle is typically $\mathcal{O}(1^\circ)$ for DB and 100 MeV masses, and less for other production modes. Since this is of the order of the angular resolution of ICARUS, it should be feasible to reconstruct the signal events as separated tracks for masses not too light.

Specific to the DUNE-MPD detector, in [14,101,102] it has been shown that boosted signal events have a narrower angular separation compared to the more isotropic background distributions. Due to this difference the search can effectively be rendered background free. Even in this case the two decaying particles must be reconstructed as separate particles, which these references claim to achieve. In these studies, the typical opening angle between the decay products is comparable to the weakly coupled scenarios we consider in this work. Therefore, we take this search to be background free.

If instead the two decay products are not separated, the event will be reconstructed as a single electron event. The background to this kind of event comes from ν_e charged-current scattering, from ν - e elastic scattering events or ν neutral-current quasielastic events. The idea of decaying

particle hiding behind the single electron signature has been explored in [103,104]. In these works it is shown how to recast the analysis of the LSND experiment (with a beam energy of 0.8 GeV) that looked for ν_e charged-current scattering [105,106] to put bounds on BSM particles decaying into e^-e^+ , hiding as single electron events. In particular, the decaying particle would present as an excess of high-energy electron events near the maximum value of the energy analyzed (200 MeV).¹¹

Considering now a higher-beam energy experiment, searches for charged-current at NuMI based experiment (with beam energy 120 GeV) typically look for neutrino with GeV energies, a bit lower than the typical LDSP energy (see for example [107]). Other scattering analysis typically look for ν - e elastic scattering. The cuts imposed require a low-energy recoil, and a very forward electron. It is unclear whether or not they can be used to put stringent bounds on misidentified e^-e^+ pairs. It will be very interesting to explore this signature of single electron hiding in the high-energy tail of scattering events at high beam-energy experiments, but we will not study this signature here.

We would like to point out that it should be possible to run the neutrino experiments in a beam dump mode, essentially removing all the background, while keeping almost all the signal (except the one coming from charged meson decays). A beam dump proposal for DUNE has been studied in [108], showing that indeed running in the beam dump mode allows neglecting all the SM backgrounds in the DUNE detector, albeit at a reduced luminosity of roughly two orders of magnitude (one order of magnitude, for the optimistic scenario). The idea to suppress neutrino background by steering the beam off the target (as in the beam-dump mode) has been already implemented at MiniBooNE [86] (although looking for DM scattering events) and MicroBooNE experiments. We will recast the MicroBooNE search, which however is quite different in spirit from the typical beam dump search, as it is optimized to look for kaon decay at rest (KDAR). The idea is to look for the decay products of kaons decaying at rest in the NuMI hadron absorber, which have a very peculiar directionality; in usual cases, the decay products of produced kaons are collected by a detector placed further down the beamline, whereas here the MicroBooNE detector is placed on the back side of the NuMI hadron absorber (e.g., see Fig. 1 in Ref. [77]). This peculiarity allows the signal events to be easily distinguished from background

¹¹We do not recast LSND bounds in the DB and DY modes, even if the intensity is one of the highest. For DY, due to a very low beam energy of 0.8 GeV, the integration range of the partonic center of mass energy is very small. For DB, the condition on the integration domain of Eq. (27) is very constraining. Relaxing the condition, by setting the rhs of Eq. (27) to 1, the bounds are still worse than DUNE.

events, with an estimated efficiency of 0.14 on signal selection.

In addition to the backgrounds discussed so far, there is an extra component coming from neutrino trident events, in which a neutrino scatters against a nucleus in a purely electroweak process, to produce a lepton-antilepton pair. As argued in [14,109], the expected number of events is $\mathcal{O}(10)$ events at DUNE-MPD, while it is lower in other liquid argon detectors like ICARUS. In more conventional detectors like NO ν A, $\mathcal{O}(10)e^-e^+$ trident events are expected [17]. Given the rather peculiar kinematics, it's possible to bring down these background events and neglect them in the analysis [14].

For all these reasons, we will compute the signal yield contours for 10 and 100 event lines when discussing prospects for DUNE-MPD. Indeed ten events represent a reasonable proxy for an almost background free search in the presence of $\mathcal{O}(1)$ experimental efficiency, although this number could be brought down in specific experiments by a more careful analysis of the reducible backgrounds. The 100 event lines instead can be representative of some signal loss due to selection cuts (which could be present for example in the weakly coupled case due to a small angular separation) or for a reduction in N_{POT} , for example due to running in dump mode for a limited amount of time.

From the experimental analysis we recast, we use 95% confidence level, including the signal efficiencies reported. For CHARM [78] which observed 0 event, we set a bound at 95% confidence level of $N_{\text{signal}} < 3$, using efficiency of 0.51 and 0.85 for the e^+e^- and $\mu^+\mu^-$ modes respectively. For MicroBooNE KDAR [77] which observed one event compared to a background expectation of 1.9 events, we require $N_{\text{signal}} < 3.8$ at 95% confidence level, with a signal reconstruction efficiency of 0.14.

V. RESULTS

In this section we present our bounds on the parameters Λ_{UV} , Λ_{IR} . As we have argued before, the dominant production mode is through the Z -portal, but the decay can proceed through either Z -portal or Higgs portal. As discussed in Sec. III, there are three production modes for DS states, each of which has a different distribution in p_{DS} and therefore contributes differently depending on the energy and the detector geometries. In Fig. 7 we show the exclusion regions individually for the three production modes, keeping to the Z portal decay for simplicity, for both weakly and strongly coupled benchmark scenarios. We will assume that the LDSP has spin-1 in order to fix the decay parametrics. In Fig. 8, we show the combined bounds from this work, for both Z and Higgs portal decays, and compare against bounds from LHC and LEP from Ref. [30]. The bounds for the Z -aligned $J_\mu^{\text{SM}}J_{\text{DS}}^\mu$ portal can be obtained from the Z portal bounds simply by rescaling the signal cross section as $\sigma_S^Z/\sigma_S^{JJ} \sim (\kappa_Z^2/\kappa_J^2)(v^2/m_Z^2)$. The case of

generic JJ portal is obtained from the Z portal by an appropriate combination vector/axial parts of the current and a rescaling of the couplings.

One can understand the general features of the excluded regions. The right edge of excluded region (on the high Λ_{IR} side) is due to the LDSPs being too short-lived. Therefore, the bound is roughly set by $\beta\gamma c\tau \sim \ell_{\text{detector}}$. Given that typically for high Λ_{IR} , n_{LDSP} for the weakly and strongly coupled cases are very similar, the right edge of the excluded region is almost identical for the weakly and strongly coupled cases. The left edge of the excluded region is instead due to the LDSPs being too long-lived, $\beta\gamma c\tau \gg \ell_{\text{detector}}$. In this case the exponentials appearing in the decay probability [Eq. (C2)] can be expanded to linear order and it is possible to get the slope of the right edge. For example, in the weakly coupled case with both production and decay through Z portal (or any generic $6D$ operator) the right edge is set by a constant ($\Lambda_{\text{IR}}^6/\Lambda_{\text{UV}}^8$) line, where the factors come from the exponential expansion to linear order and the $1/\Lambda_{\text{UV}}^4$ in the production cross section. For the strongly coupled case the n_{LDSP}^2 factors have a nontrivial dependence on Λ_{IR} , and in general it is not possible to get the slope. Its effect is to increase the total signal events since the number of LDSPs is larger than in the weakly coupled scenario considered. This enhancement asymptotes to the weakly coupled case when Λ_{IR} is close to threshold due to the behavior of the chosen function.

In all these exclusion plots, regions bounded by solid lines show the excluded parameter space from recasts of past and recent DS searches (CHARM [78] and MicroBooNE [77]). The region bounded by dashed contours in our plots show the potential of current and future upcoming neutrino experiments; NO ν A-ND and ICARUS (current), and DUNE-MPD (near future). To compare their potential with future DS experiments, we also show the projections coming from the future beam dump experiment SHiP. For these projected exclusions, we have shown the ten signal events line assuming 100% reconstruction and detection efficiency. Following our discussion in Sec. IV, for DUNE-MPD, we also show the 100 events line in Fig. 8.

There are several features of the bounds which make the neutrino experiments a very powerful probe for dark sectors, in the parametrization considered in this work. First of all, we find that the bounds from current and upcoming neutrino experiments are comparable to dedicated DS experiments, with a reach of Λ_{UV} in multiple TeV range, for Λ_{IR} in the MeV–GeV range. This is similar to the ranges probed in high-energy experiments like LHC and LEP (as done in Ref. [30]). The typical scale l at which the detectors are placed is much larger than the corresponding scale in the LHC DV searches, and the typical boosts involved are also different, which together select a somewhat larger τ and hence a smaller Λ_{IR} region compared to LHC. Importantly, neutrino experiments fill the gaps in the parameter space

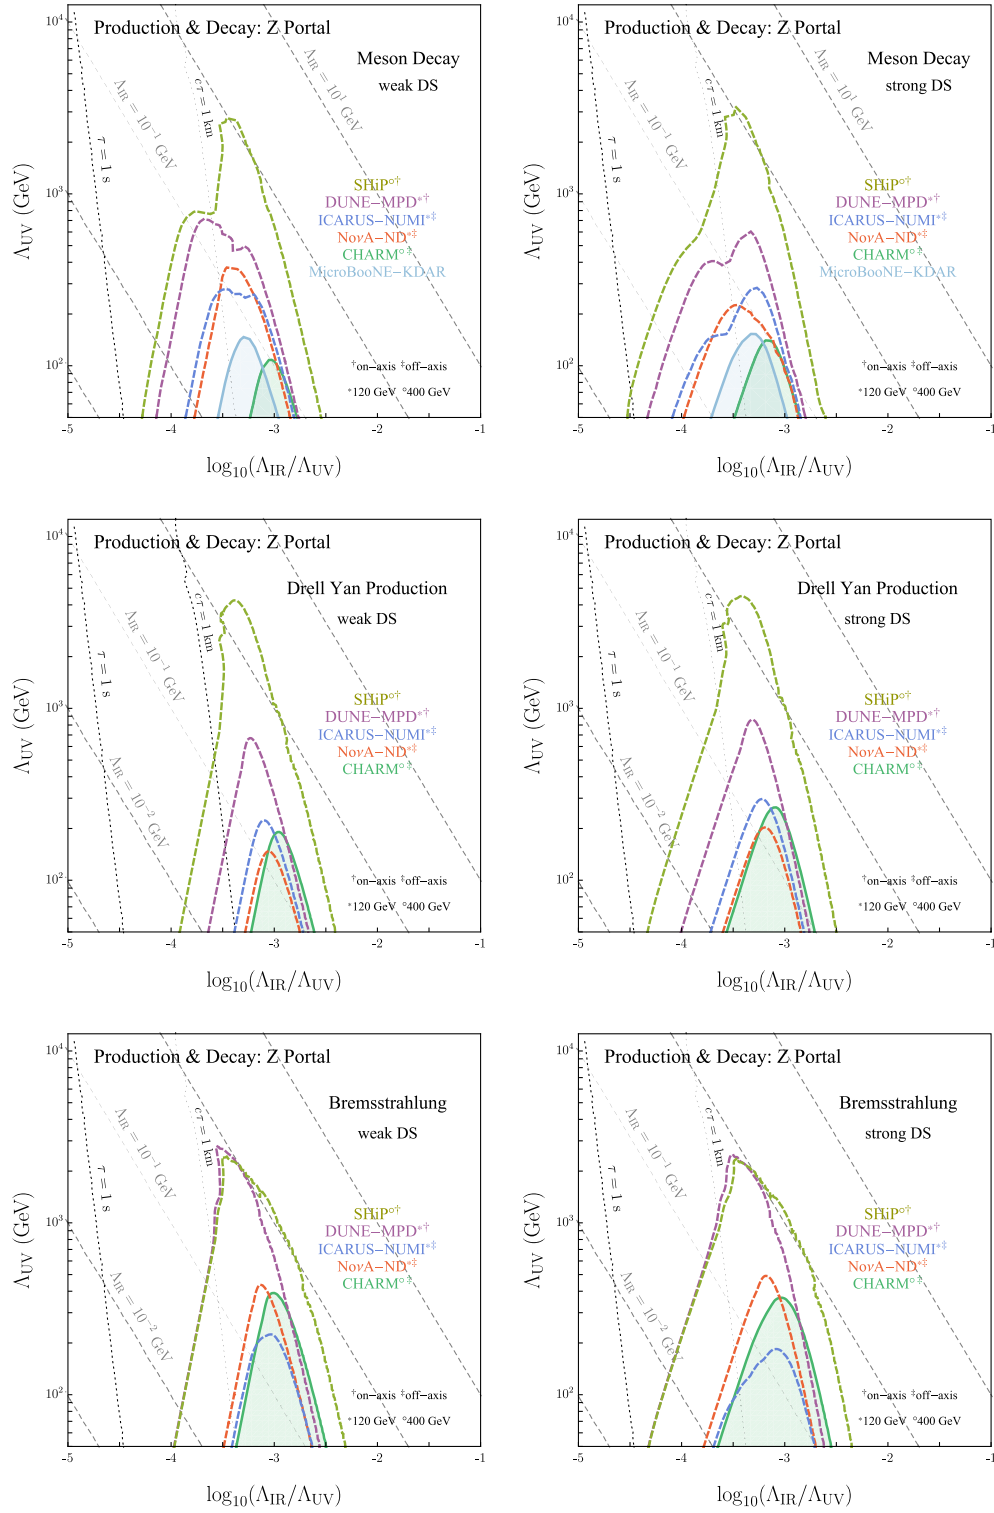


FIG. 7. Constraints on DS production and decay from $D = 6$ Z portal for various production modes: Meson decay (top row), Drell-Yan (mid row) and dark bremsstrahlung (bottom row), for weakly coupled (left) and strongly coupled (right) benchmark cases. Exclusions are shown in solid lines, while future projections are shown in dotted lines. The superscripts indicate the beam energy and the on/off axis nature of detectors for each of the experiments considered. The black dotted lines show the LDSP lifetime τ isocurves, while the gray dashed lines show Λ_{IR} isocurves. The bounds assume $p_{DS}^2/\Lambda_{UV}^2 < 0.1$ for EFT validity which is satisfied by restricting to $\Lambda_{UV} > 50$ GeV. All plots assume $\kappa_i^2 c_i = 1$, where i labels the portal, κ is the portal coupling and c is a measure of degrees of freedom of the DS.

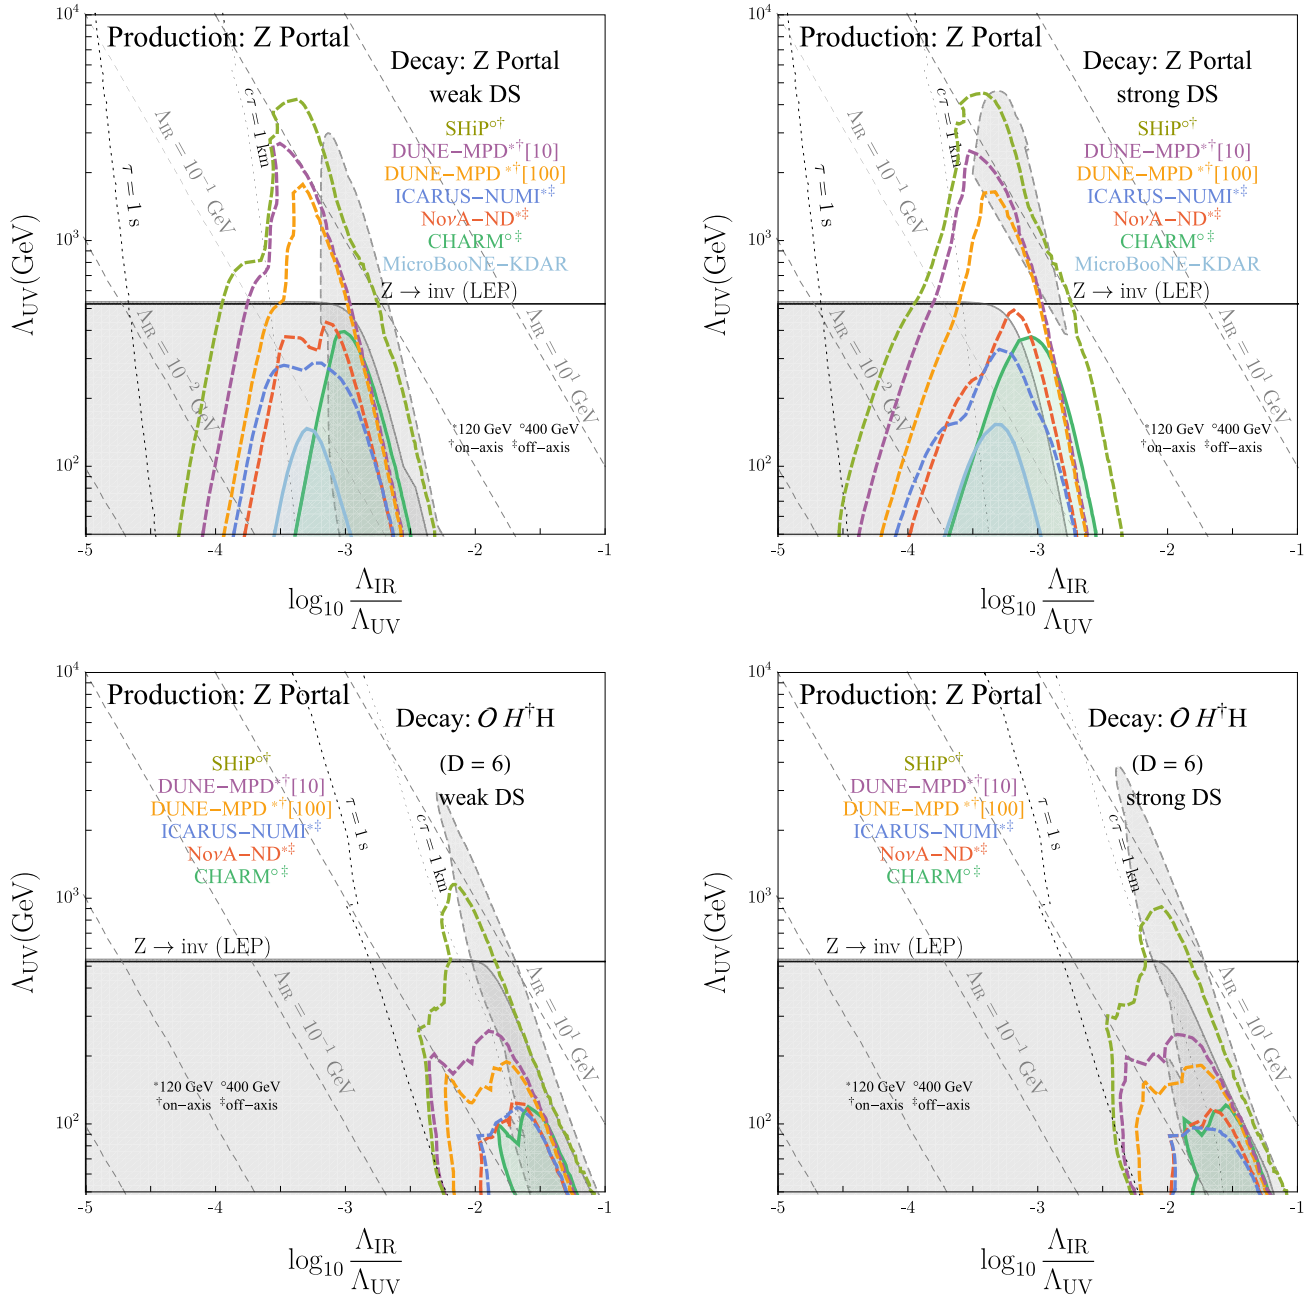


FIG. 8. Constraints on DS production through the $D = 6$ Z portal, and decay through the same Z portal (top), or through $D = 6$ $\mathcal{O}H^+H$ portal (bottom), at various neutrino experiments. We have shown both the 10 event and the 100 event lines for DUNE. For comparison, bounds from high-energy colliders (obtained in Ref. [30]) are also shown in gray (monojet searches with solid boundary and displaced vertex searches with dashed boundary). The left (right) plots assume weakly coupled (strongly coupled) dark dynamics. The exclusion from the Z invisible width measurement at LEP is shown by the horizontal solid black line. We restrict to $\Lambda_{UV} > 50$ GeV for EFT validity. All plots assume $\kappa_i^2 c_i = 1$, where i labels the portal, κ is the portal coupling and c is a measure of degrees of freedom of the DS.

coming from trigger¹² and event selection requirements at LHC and LEP, since they are sensitive to much lower energy

¹²This could be improved with dedicated trigger designs, as done in [110] for CMS in a different lifetime region than the typical one of the gap.

activity in the detector. Even more importantly, in portals which are not enhanced by a resonant production, the EFT condition $(p_{DS}^2)_{\max} < \Lambda_{UV}^2$ makes LHC and LEP bounds inconsistent, an issue which is again alleviated at neutrino experiments due to a smaller \sqrt{s} involved. All these features are seen in the bounds in Figs. 7 and 8.

In the following subsections, we will discuss in detail how prospective DS searches at neutrino experiments can complement current bounds for different portals considered in this work. We will also emphasize the difference between the production modes, especially on how a particular detector geometry can favor one mode over the other.

A. Z-portal production

Consider first the MD production mode through the Z portal (Fig. 7 first row), where we show bounds from various experiments. We also show bounds coming from the MicroBooNE KDAR analysis [77] which is only relevant for the MD mode. We find that for neutrino experiments based at 120 GeV proton beam, in general the strongest bounds come from radiative decays of K meson ($K \rightarrow \pi + DS$) due to the large number of K mesons produced with respect to other mesons. However, for the strongly coupled DS case where the kinematic condition on p_{DS}^2 is stronger due to a larger n_{LDSP} , we find that at DUNE and ICARUS, $\phi \rightarrow DS$ decays can give a stronger bound on Λ_{UV} as compared to K meson bounds. For experiments based on 400 GeV proton beam, heavier mesons like $B, J/\psi$ can be produced in large numbers and can contribute to bounds at CHARM and SHiP. These heavier mesons can in principle probe larger Λ_{IR} due to the relaxed kinematic condition $n_{LDSP}\Lambda_{IR} \lesssim M$, where M is the mass of the decaying meson. The Λ_{IR} reach is correlated with the Λ_{UV} reach, for fixed lifetime. At SHiP, we indeed find that J/ψ decays probe the highest Λ_{UV} scales as opposed to K -meson decays which suffer from the kinematic condition. However, we find that K mesons can still improve reach on lower Λ_{IR} values relative to $J/\psi, B$ mesons due to larger geometric acceptance for $K \rightarrow \pi + DS$ and larger number of K mesons. At CHARM, we find that $J/\psi \rightarrow DS$ decays give the leading bounds which dominate those coming from $K \rightarrow \pi + DS$ decays.

For the Drell-Yan (direct partonic) production mode (Fig. 7 middle row), the best bounds come from SHiP, and the other experiments only give subleading bounds. Within them, due to the collinear nature of the produced DS beam, detectors of the on-axis type are more sensitive to this mode. Note that compared to the DB mode, the average DS boost is smaller for the DY mode, so that the LDSP spread is more and the off-axis detectors are penalized less.

For the case of dark bremsstrahlung (Fig. 7 bottom row), the DS is produced very collimated along the beam line, favoring detector geometries closer to it. We find that the best bounds from neutrino experiments for this mode come from DUNE-MPD. These are comparable in Λ_{UV} and only probe slightly smaller Λ_{IR} values, as compared to the future beam-dump experiment SHiP. This is because in bremsstrahlung the typical p_{DS}^2 is cut roughly around QCD scales, and an increase in \sqrt{s} at SHiP as compared to DUNE does not lead to a large increase in the production cross section. The other experiments shown, ICARUS-NUMI, NO ν A and

CHARM give subleading reach in both Λ_{UV} and Λ_{IR} . Despite an increase in the number of POTs with respect to NO ν A, ICARUS still has lower sensitivity due to a reduced angular coverage. Both ICARUS and CHARM, due to their off-axis nature, miss out signal events from the forward DS beam, characteristic of the dark bremsstrahlung mode.

Combining all the production modes, in Fig. 8 we show the final excluded parameter space, for both weakly and strongly coupled benchmarks. We also show the results from [30] which studied resonant DS production through Z portals at high-energy colliders and presented exclusion regions from ATLAS monojet search [111], displaced vertex search [112,113], and from total Z-width bounds from LEP [114]. The bounds presented here probe different parts of the parameter space, in particular in Λ_{IR} , even if the Λ_{UV} reach is comparable to before, and also probe gaps in parameter space in earlier work that came from trigger requirements. The complementarity of the bounds at neutrino experiments, as compared to missing energy and displaced searches at the LHC is due to the peculiar position of the near detectors of neutrino experiments, placed at $\mathcal{O}(10^2)$ meters. We also find that the bounds are stronger than past beam dump searches like E137 and NA64 (whose results can be found in [30,33]) due to a larger N_{POT} , and in some cases, larger beam energy.

From Fig. 8 we see that one of the current strongest bounds for Z portal DS still comes from the indirect Z width measurement at the LEP, $\Lambda_{UV} > 525(k_J^2 c_J)^{1/4}$ GeV [30] which is independent of Λ_{IR} till the kinematic threshold. This bound is stronger than CHARM and MicroBooNE. Additionally, CHARM and MicroBooNE are also weaker than the LHC monojet and displaced vertex bounds except for the strongly coupled case (Fig. 8 top right) where they probe slightly higher Λ_{IR} values. Prospective DS searches at current Fermilab neutrino facilities, ICARUS and NO ν A-ND, improve on CHARM and MicroBooNE, but they are still weaker than the LEP bound.

Most importantly however, we find that future neutrino detector DUNE-MPD will be sensitive to Λ_{IR} in the range $\mathcal{O}(0.1 - 1)$ GeV for Λ_{UV} of few TeVs, a region not covered by LHC exclusions. Future LLP experiment SHiP based on 400 GeV proton beam would further improve sensitivity with respect to DUNE-MPD and LHC searches. These improvements are either due to a higher geometric acceptance from being on-axis or from having a wider detector, or due to a higher beam energy.

Another past proton beam dump experiment ν CAL, with a beam energy of 70 GeV, has searched for the decay of a scalar particle [115,116]. Recasting this search [117,118] gives bounds slightly better than CHARM, excluding in roughly the same Λ_{IR} region, and with Λ_{UV} up to 1 TeV. This is due to the fact that the bremsstrahlung cross section for this energy is the same as of CHARM, given the proton virtuality cutoff, the experiment has a similar luminosity, but it is closer to the beam and target. The bounds are still

weaker than DUNE and SHiP. Recasting the Heavy Neutral Lepton search of PS191 [119,120], as suggested in [121,122], gives bounds better than CHARM but worse than ν CAL and DUNE. We do not show these bounds to keep the plots uncluttered.

Emphasizing the complementarity of neutrino experiments with respect to the LHC searches, we note that for the strongly coupled case (Fig. 8 top right), the LHC (ATLAS) searches are not sensitive in a gap of parameter space values close to $\Lambda_{\text{UV}} \sim 500 \text{ GeV} - 2 \text{ TeV}$ and $\Lambda_{\text{IR}} \sim 0.1-0.5 \text{ GeV}$. We find that both future neutrino experiment DUNE-MPD and dedicated LLP experiment SHiP would remarkably fill this gap in the Z portal DS parameter space. These gaps were due to trigger and event selection requirements.

So far we have only considered decay through the Z -portal, but the decay can also proceed through the Higgs portal. Before proceeding with that, a couple of comments about the interplay of the quantum numbers of LDSPs and the relevant decay portals is in order. According to the Landau-Yang theorem [123,124], a massive spin-1 particle cannot decay into two massless spin-1 particles. This implies that if the LDSP is a spin-1 particle then for values of $\Lambda_{\text{IR}} < 2m_e$, the only decay channel is to three photons, the decay width for which is very small, $\Gamma \sim \alpha_{\text{EM}}^3 \Lambda_{\text{IR}} (\Lambda_{\text{IR}}/\Lambda_{\text{UV}})^4 (\Lambda_{\text{IR}}/m_e)^8$, and effectively the only signal is a missing energy; the plots are not changed in the shown range. If instead the LDSP is a spin-0 particle, there is no such condition. However, for such small values of Λ_{IR} , the LDSPs are too long-lived and cannot be efficiently constrained at the experiments considered here. For this reason it is crucial to realize that while production from the Higgs portal is very suppressed, it might be relevant for decays if it is the only available decay mode. We should clarify that if multiple portals are available for decay, one has to consider the dominant one. When considering the LDSP decay through the Higgs portal, we are assuming it to be dominant compared to other portals. Note that a spin-1 LDSP cannot decay through the Higgs portal due to quantum numbers.

In the bottom row of Fig. 8 we show the bounds where the LDSP decay occurs through the $\Delta = 4$ Higgs portal. The longer lifetime of the LDSPs, due to a very small coupling to leptons and the extra $\Lambda_{\text{IR}}^2/m_h^2$ suppression factor, effectively shifts the exclusion regions to higher Λ_{IR} regions. However, bounds from such values of Λ_{IR} can be suppressed due to being too close to the edge of allowed phase space, effectively chopping off the bounded region. This makes these bounds typically weaker than colliders, in their Λ_{UV} reach, although they still cover regions unconstrained by Z portal decays at higher $(\Lambda_{\text{IR}}/\Lambda_{\text{UV}})$ ratio.

Finally, for completeness, we will now tabulate constraints coming from invisible meson decays where the LDSP is long-lived enough to escape detectors. Overall, we find that these bounds are weaker in their Λ_{UV} reach than the ones coming from both LHC and neutrino detectors.

For the Z portal (both production and decay), the strongest constraints from invisible meson decay come from flavor-changing decays of B and K mesons (updated with respect to Ref. [30]). We take the *BABAR* upper limit for $B^+ \rightarrow K^+ \bar{\nu}\nu$ decays [125]: $\mathcal{B}(B^+ \rightarrow K^+ \bar{\nu}\nu) < 1.6 \times 10^{-5}$, which gives

$$\frac{\Lambda_{\text{UV}}}{\text{GeV}} > 60(\kappa_J^2 c_J)^{1/4},$$

$$\frac{\Lambda_{\text{IR}}}{\text{MeV}} \ll 108(\kappa_J^2 c_J)^{-0.2}(\text{weak}), 65(\kappa_J^2 c_J)^{-0.2}(\text{strong}). \quad (36)$$

For the case of $K \rightarrow \pi + \text{DS}$, we take the upper limit from the NA62 Collaboration [126]; $\mathcal{B}(K^+ \rightarrow \pi^+ + \bar{\nu}\nu) < 1.06 \times 10^{-10}$ which gives

$$\frac{\Lambda_{\text{UV}}}{\text{GeV}} > 68.8(\kappa_J^2 c_J)^{1/4},$$

$$\frac{\Lambda_{\text{IR}}}{\text{MeV}} \ll 83(\kappa_J^2 c_J)^{-0.2}(\text{strong}) \quad (37)$$

The bounds for a weakly coupled DS are similar. For the case in which the LDSP decay is via $\Delta_{\mathcal{O}} = 4$ Higgs portal instead, the bounds for the B - and K -meson decays respectively are

$$\frac{\Lambda_{\text{UV}}}{\text{GeV}} > 60(\kappa_J^2 c_J)^{1/4}, \quad \text{for } \frac{\Lambda_{\text{IR}}}{\text{GeV}} \ll 1.6(\kappa_J^2 c_J)^{-0.14}, \quad (38)$$

$$\frac{\Lambda_{\text{UV}}}{\text{GeV}} > 68.8(\kappa_J^2 c_J)^{1/4}, \quad \text{for } \frac{\Lambda_{\text{IR}}}{\text{MeV}} \ll 77(\kappa_J^2 c_J)^{-0.14}. \quad (39)$$

In the above, the condition on Λ_{IR} has been calculated assuming a strongly coupled DS, and they do not change significantly for the weakly coupled case.

The invisible decays from J/ψ which have been searched for by the BES Collaboration [127] set an upper limit on $\mathcal{B}(J/\psi \rightarrow \bar{\nu}\nu) < 7.2 \times 10^{-4}$. However, we found the resulting bound on Λ_{UV} to be weaker than those coming from the *BABAR* and NA62 limits on B , K decays, and we do not report it here.

A preliminary study of FASER and SND, using an integrated luminosity of $\mathcal{L} = 150 \text{ fb}^{-1}$ and three signal event exclusion, shows that the exclusion power of these experiment comes mostly from the decays of SM forward object like mesons or on shell Z , whose spectra can be taken from the FORESEE package [128]. For the decay through Z portal, the bounds are comparable to DUNE but are in a slightly different Λ_{IR} region due to a different boost, and a different distance at which the detectors are located, while for decays through Higgs portal ($\Delta = 4$) they can improve the bounds up to $\Lambda_{\text{UV}} 1 \text{ TeV}$; this is because for resonant production through Z portal, $p_{\text{DS}}^2 \simeq m_z^2$, allowing for larger LDSP masses to be produced and tested. Future experiments like FASER2 can exclude up to 10 TeV given the larger luminosity and the dimensions of the detector,

compared to its predecessor. These will be studied in detail in a future work, together with other experiments at the lifetime frontier such as MATHUSLA and Codex-b.

B. JJ portal (Z -aligned) production

Even though the Z -aligned JJ portal and Z -portal are equivalent at neutrino experiments after an appropriate rescaling of the κ , there is a distinction between them at high-energy experiments that can produce a Z on shell. Contrary to Z -portal case, for the Z -aligned JJ portal, the LHC bounds are generally weaker due to the lack of resonant production and EFT consistency condition on Λ_{UV} . For the same reason, there is no bound coming from Z -width. In this scenario, the bounds on Λ_{UV} come only from LEP missing-energy searches (see Fig. 8 in Ref. [30]). Therefore, regions in the parameter space with too short lifetimes are not tested due to the requirement for the LDSP to decay outside the detector. On the other hand, the bounds coming from high-intensity experiments such as neutrino experiments are essentially unchanged with respect to the Z -portal case, so that all the discussion from before applies; they are able to exclude a larger portion of Λ_{UV} by roughly one order of magnitude in the large lifetime region (low Λ_{IR}), compared to the LEP/LHC detector size, while it excludes a completely unexplored region at small lifetime (or large Λ_{IR}).

C. JJ portal (generic) production

The previous sections can give us an insight on how high-intensity experiments can put a bound on a generic $J_\mu^{SM} J_{DS}^\mu$ portal, where J_μ^{SM} is a generic flavor-conserving SM current. Missing-energy bounds coming from LHC will still hold provided $\mathcal{O}(1)$ couplings to light quarks, although as explained in the previous section they are limited by the EFT condition. If these are absent, (e.g., for $\bar{e}\gamma^\mu e J_\mu^{DS}$ portal) the leading bounds come from LEP mono-photon searches ($\Lambda_{UV} \gtrsim 10^2$ GeV for $\Lambda_{IR} \lesssim 100$ MeV), while electron beam dump experiments like E137 and missing energy searches at NA64 put weaker bounds, see Ref. [30].

Proton beam-based neutrino experiments cannot probe hadrophobic current interactions given that couplings to quarks are essential for all production modes. Since neutrino experiments typically exclude LDSP masses for $\Lambda_{IR} \lesssim 2m_\pi$, if the decay proceeds through generic JJ portal, we need nonzero couplings to electrons. If that is small, the Higgs portal may be relevant depending on couplings. This feature is not present in missing-energy searches at high-energy colliders and high-intensity experiments, which only probe the production mode. This problem can be circumvented if instead of looking at displaced vertex signatures (where LDSP decays inside the detector), in which both DS production and DS decay into SM are required, scattering events are also considered.

As mentioned in previous sections, we do not look at such signatures due to the extra assumptions needed with respect to LDSP decays.

We remark that no big difference is expected from changing the axial or vector nature of the SM current as long as their coupling is of the same order. While for Z portal the axial contribution to bremsstrahlung is larger than the vector counterpart, due to the accidentally small coupling of the vector component, this is not necessary for a generic case. A similar argument also hold for DY mode, while for MD mode the quantum numbers of the SM current select the relevant meson processes (see Appendix B for details). To conclude, as long as DS has a coupling to proton and electrons in J_{SM}^μ , we expect the results to not change dramatically at fixed magnitude of the couplings; the bounds presented in Sec. VA apply.

D. Higgs-portal production

The bounds at neutrino experiment for production through $\Delta_\mathcal{O} = 4$ Higgs portal $\mathcal{O}H^\dagger H$ are very weak; assuming decay through Z portal $\Lambda_{UV} \ll 10^2$ GeV for DB and DY modes since these modes are suppressed by a small Higgs coupling. Only radiative meson decays happening through a top loop do not suffer from such a problem. The strongest bounds for this case then come from meson decay $K \rightarrow DS$ where for DUNE-MPD, we get $\Lambda_{UV} \lesssim 12$ GeV. This exclusion is much weaker than the bounds coming from missing-energy searches at LHC and Higgs coupling fits [30], $\Lambda_{UV} \gtrsim 450$ GeV. For $\Delta \geq 4$ the rate is suppressed with respect to the Z portal as explained in Sec. III A. The situation is slightly improved for high-energy beam experiments like SHiP (where $\Lambda_{UV} \lesssim 74$ GeV is excluded for the $B \rightarrow K + DS$ meson decay), but is still not competitive with the ones coming from Higgs resonant production at LHC. For this reason, we do not show any plots for production through the Higgs portal.

For a $\Delta_\mathcal{O} = 3$ Higgs portal, at DUNE-MPD, $K \rightarrow DS$ decay gives the leading bound, $\Lambda_{UV} \lesssim 540$ GeV. At SHiP, we find that the leading bounds come from $B \rightarrow K + DS$ decays which give $\Lambda_{UV} \lesssim 10$ TeV for $\Lambda_{IR} \sim 1.8\text{--}2.8$ GeV. The DUNE-MPD bounds are weaker than the LHC missing-energy searches (which exclude $\Lambda_{UV} \lesssim 8$ TeV) whereas the SHiP bounds are stronger.

The forward experiment FASER does not put strong constraint on DS production and decay through resonant Higgs portal. Only FASER2, due to a higher luminosity and geometric acceptance, can put bounds that are comparable with the conventional LHC searches in a narrow lifetime regions; up to $\Lambda_{UV} = 10$ TeV for $\Delta = 3$, while up to $\Lambda_{UV} = 1$ TeV for $\Delta = 4$.

VI. SUMMARY AND DISCUSSION

Secluded sectors that interact very feebly with the SM have the potential to be probed at the high-intensity frontier,

particularly at neutrino experiments (as has been previously explored in Refs. [14,16,51,129], see also Refs. [117,118]). Most of the past work has focused on the case of relevant portals, while the case of irrelevant portal DS scenario has only recently been explored [30,32,33,102]. In this work, we have considered the sensitivity of DS that interacts with SM through a dimension-6 irrelevant portal, at past and current neutrino experiments, and its prospective discovery in both existing and future neutrino experiments based on proton beams.

We have performed a detailed study of the possible production mechanisms of DS through nonrenormalizable portals; meson decays ($M \rightarrow m + DS$, $V \rightarrow DS$), direct partonic production ($\bar{q}q \rightarrow DS$, $gg \rightarrow DS$), and dark bremsstrahlung ($pp \rightarrow DS + X$). The interplay between the various production mechanisms as a function of the DS invariant mass squared p_{DS}^2 can be summarized in the plot shown in Fig. 1. Compared to previous works on irrelevant portals, we have added production details, and also considered strongly coupled dark sectors, and done so in a model-agnostic framework. Further, we have constrained such dark sectors using past and current analyses at beam dump/neutrino experiments, also showing projections for prospective searches at existing and future neutrino experiments. In order to emphasize the importance of these bounds, we have also compared our results with previous bounds on such portals.

In an earlier work of this scenario [30], the most stringent bounds on DS excitations produced from the decay of Z bosons was set by LHC monojet searches [111] and LHC displaced vertex search [112,113], in a range of (Λ_{UV} , Λ_{IR}) values dictated by various factors such as the energy of the experiment and the lifetime of the DS etc. In the present work, we have tried to address the question if neutrino experiments, being placed farther from the interaction point (as compared to say the ATLAS detector at LHC) could probe a lower Λ_{IR} range, thereby testing a complementary parameter space with respect to high-energy colliders for such elusive dark sectors. Our main summary plots can be found in Fig. 8.

While the present work focuses on the utility of neutrino experiments for probing dark sectors, we would like to mention the status of other probes of the dark sectors considered here, for completeness. Colliders and beam dump probes produce the DS states directly. Other setups that also produce DS states directly result in astrophysical bounds from supernova cooling, from lifetime of horizontal branch stars and from positronium lifetime. Due to kinematics, bounds coming from astrophysical objects such as supernovae cannot probe the $\Lambda_{IR} \gtrsim 100$ MeV given the lower typical temperature. Therefore, they are subleading in the region in which neutrino experiments are competitive with respect to LHC bounds, and we will not show them. However in the much lower Λ_{IR} regime, they can become the most competitive bounds, as can be seen for example by recasting the results of [130]. Complementary to those are

indirect probes where the initial and final states are SM states, and DS degrees of freedom propagate internally. Examples of such probes are electroweak precision tests (EWPT), fifth-force constraints, torsion balance experiments, molecular spectroscopy, etc. Depending on the process, these indirect probes are UV sensitive (and in that case they do not probe the dark dynamics directly) or give weaker constraints. A careful analysis of all these direct and indirect effects was already carried out in Ref. [30] and we refer the reader to there.

In this work, for the case of Z portal DS production, we find that past analyses and prospective DS searches at current neutrino experiments give weaker bounds when compared with the current bounds from LHC and LEP in resonant production scenarios. However, future neutrino experiments such as DUNE-MPD would improve on this, and will be sensitive to Λ_{IR} in the range 0.1–1 GeV for $\Lambda_{UV} \sim 1$ TeV. The current displaced vertex searches at LHC are already probing Λ_{UV} as high as few TeVs, but only in the Λ_{IR} range ~ 0.6 – 2.5 GeV for the strongly coupled DS case. The ATLAS DV searches lose sensitivity in the range of $\Lambda_{IR} \sim 0.1$ – 0.6 due to trigger requirements (as can be seen from the gap on the right plot in Fig. 8).¹³ Future neutrino experiment DUNE-MPD will have a unique sensitivity to access this gap in the parameter space for a range of $\Lambda_{IR} \sim 0.1$ – 1 GeV for Λ_{UV} of a few TeVs.

We have also compared these results with projections from the proposed experiment SHiP which serves as a benchmark for LLP experiments. As can be seen from Fig. 8, SHiP will improve on the reach of DUNE-MPD by probing Λ_{UV} of few TeVs for a range of $\Lambda_{IR} \sim 0.1$ – 2 GeV. This is mainly due to its higher-proton beam energy of 400 GeV and larger geometric acceptance ($\epsilon_{geo} \sim 1$ for the production modes bremsstrahlung and DY and $\epsilon_{geo} \sim 0.1$ – 0.9 for K and J/ψ decays).

As we have described in the previous section, we can recycle our bounds at neutrino experiments on Z portal also for the case of $J_{\mu}^{SM} J_{DS}^{\mu}$ portal where $J_{\mu}^{SM} = \bar{f}\gamma_{\mu}f, \bar{f}\gamma_{\mu}\gamma^5 f$, $f = l, q$. The earlier work in [30] found no bounds from LHC for this portal, once EFT considerations were taken into account. The only constraint presented is the one for $J_{\mu}^{SM} = \bar{e}\gamma_{\mu}e$ from monophoton searches at LEP (see Fig. 8 in Ref. [30]) where the excluded space is restricted to $\Lambda_{UV} \lesssim 200$ GeV for $\Lambda_{IR} \lesssim 0.1$ GeV. These bounds are much weaker than the bounds we get at neutrino experiments. Therefore neutrino experiments are a useful tool to study DS that do not directly mix with the Z and that are not enhanced by resonant production at colliders. We have already

¹³The trigger requirements imposed in [30] depend on the n_{LDSP} distribution. Events where the number of LDSPs produced has a downward fluctuation can loosen the cut, but will also affect the total cross section, the decay probability and the geometric efficiency. Including this effect systematically will reduce the unprobed region but not entirely.

explained how our Z portal bounds can be recycled for a JJ portal, since a JJ portal can be obtained from Z portal after integrating out Z mediator.

For the case of the Higgs portal $\mathcal{O}H^\dagger H$ production (for $\Delta_{\mathcal{O}} = 4$) the bounds from neutrino experiments are very weak, and are limited to values of Λ_{UV} much below the electroweak scale. Whereas, in comparison, bounds from Higgs resonant production derived in Ref. [30] coming from missing energy and displaced vertex searches at high-energy colliders are much stronger.

The bounds presented here are derived under a model agnostic approach, and are applicable to a large class of DS models (see Ref. [30] for explicit examples). Knowledge of the underlying dark dynamics can be used to study other possible signatures like DS scatterings with SM particles, but will need to be done on a case-by-case basis, and hence is out of the scope of this work. We have discussed in detail our assumptions and limitations of our approach in Sec. II. Our results are conservative and can be improved if the full theory is defined explicitly. Despite this, we claim that our approach can be very useful in giving a qualitative picture.

The point of this work is to convey the usefulness of a model agnostic approach to exploring dark sectors, and the potential of neutrino experiments (both current and future) as unique probes of irrelevant DS-SM portals. Future proposed LLP experiments at LHC interaction points like MATHUSLA, CODEX-b, ANUBIS are designed to improve reach on Λ_{IR} scales for such elusive DS. However, future neutrino experiment DUNE multi-purpose detector (MPD) [82] running at the LBNF (Long Baseline Neutrino Facility) would probe low Λ_{IR} scales in a shorter timescale. The forward LHC detectors like FASER and SND (see Ref. [131] for a recent status report), built for searching feebly interacting particles would be taking data during the LHC Run 3, and could also give useful bounds for our DS. We hope our study would motivate analyses of neutrino-detector data for the search of such elusive dark sectors.

ACKNOWLEDGMENTS

We would like to thank R. Contino for initial discussions and for useful input at various stages of this project. We would additionally like to thank P. Asadi, B. Batell, D. Buttazzo, N. Foppiani, S. Homiller, W. Jang, K. Max, F. Poppi, D. Redigolo, M. Reece, M. Scodeggio, M. Strassler, and L. Vittorio for reading the draft, for useful discussions, and for pointing us to relevant references. We also thank Adam Ritz and Saeid Foroughi-Abari for clarifications on their work on proton bremsstrahlung. M. C. is partially supported by Progetti di Ricerca di Interesse Nazionale (PRIN) 2017L5W2PT. R. K. M. is supported by the National Science Foundation under Grants No. NSF PHY-1748958 and No. NSF PHY-1915071. S. V. is supported in part by the MIUR under Contract No. 2017FMJFMW (PRIN2017).

APPENDIX A: TIMELIKE FORM FACTORS

In the extended vector-meson dominance (eVMD) formalism, form factors are modeled as sum over meson states \mathbf{m} sharing the same quantum numbers of the SM operator X involved in the matrix element. Typically the formalism is applied in the spacelike region, and the form factors are essentially the sum of the propagators of the virtual mesons. However given that the exchanged momentum p_{DS} in our case is timelike, the form factors are modeled as Breit-Wigners, to allow the virtual exchanged meson to go on shell and resonantly mix with the DS system,

$$F_{\mathbf{m}}(p_{DS}^2) = \sum_i f_{\mathbf{m}_i} \frac{m_{\mathbf{m}_i}^2}{p_{DS}^2 - m_{\mathbf{m}_i}^2 + im_{\mathbf{m}_i}\Gamma_{\mathbf{m}_i}}, \quad (\text{A1})$$

where $m_{\mathbf{m}}$, $\Gamma_{\mathbf{m}}$, and $f_{\mathbf{m}}$ are the mass, the decay width and the ‘‘couplings’’ for the meson \mathbf{m} respectively. The number of resonances in the meson tower sharing the \mathbf{m} quantum numbers is such that it allows to enforce the correct asymptotic behavior in q^2 of the form factor coming from sum rules. The couplings $f_{\mathbf{m}_i}$ are fitted from data and by the overall coupling normalization. In order to get the total form factor F_X , we need to also take into account the cutoff for too high virtuality of Eq. (28),

$$F_X = F_D(Q^2)F_{\mathbf{m}}(p_{DS}^2). \quad (\text{A2})$$

We next discuss the Higgs and the Z -portal case separately.

1. Higgs portal

The Higgs coupling to the protons is of the form

$$g_{hNN}F_H(q^2)h\bar{u}_p u_p, \quad (\text{A3})$$

where u_p is the proton spinor. The form factor F_H also includes the virtuality cutoff of Eq. (27),

$$F_H = F_D F_S. \quad (\text{A4})$$

The form factor $F_S(p_{DS}^2)$ is estimated using Eq. (A1), and we include the first three CP -even, scalar resonances. The values used for masses, width and f used in the FF are given in Table III.

2. Z portal

The effective Z vertex for the proton is modeled as

$$\begin{aligned} & \frac{g_{EW}}{\cos\theta_W}\bar{u}_p\gamma^\mu\left(F^p(q^2)\left(\frac{1}{2}-\sin^2\theta_W\right)-\sin^2\theta_W F^\omega(q^2)\right)u_p \\ & + \frac{g_{EW}}{4\cos\theta_W}g_A F_A(q^2)\bar{u}_p\gamma^\mu\gamma^5 u_p, \end{aligned} \quad (\text{A5})$$

TABLE III. Masses m_m , width Γ_m and coupling f_m of the mesons m for the scalar, vector isosinglet, vector isotriplet, and axial-isotriplet proton form factors used in the eVMD approach.

	f_0 (Scalar)			ω (Z, vector isosinglet)			ρ (Z, vector isotriplet)			a_1 (Z, axial isotriplet)	
m_m (GeV)	0.5	0.980	1.37	0.782	1.42	1.67	0.775	1.45	1.72	1.23	1.647
Γ_m (GeV)	0.275	0.5	0.35	8×10^{-3}	0.2	0.3	0.149	0.4	0.25	0.4	0.254
f_m	0.28	1.8	-0.99	1.011	-0.881	0.369	0.616	0.223	-0.339	2.26	-1.26

where q is the exchanged momentum between the virtual and real proton (p_{DS} in our case) and u_p is the proton spinor. The prefactors in front of the isosinglet- and isotriplet-vector form factors (F^ω and F^ρ , respectively) come from the decomposition of the vector piece of the quark Z -current in the singlet and triplet component in isospin space, using the approximate isospin symmetry of proton and neutron. The axial-vector form factor has a slightly different normalization than the vector one. In particular, there is an extra $g_A \simeq 1.2$ multiplying the overall form factor in Eq. (A1), as shown in Eq. (A5). We can see that the vector part is subleading with respect to the axial current due to the absence of $\sin^2 \theta_W$ suppressing factor and a $\mathcal{O}(1)$ coupling g_A . Notice that in principle, along with the Dirac-like form factors appearing in Eq. (A5), there should be the nonrenormalizable Pauli-like terms [132],

$$G_V(q^2) i \frac{[\gamma^\mu, \gamma^\nu]}{m_p} q_\nu + G_A(q^2) \gamma^5 \frac{1}{m_p} q^\mu, \quad (\text{A6})$$

where the second piece (the axial one) is mediated by pion exchange. It turns out that this contribution vanishes when contracted with the JJ correlator, due to current conservation. The vector piece, corresponding to the anomalous magnetic-proton contribution, does not vanish. Given that it is numerically subleading, and that there are large uncertainties to extrapolate such form factor in the timelike region [133,134], we will neglect this contribution.

The form factors appearing in Eq. (A5) are again computed in the eVMD formalism, as a sum of Breit-Wigners of ρ , ω , a_1 meson for F^ρ , F^ω , F_A , respectively. The coefficients used in the form factors are found in Table III.

The final form factors appearing in Eqs. (32)–(33) are defined respectively as

$$F_Z^V = F_D \left(-\sin^2 \theta_W F^\omega + F^\rho \left(\frac{1}{2} - \sin^2 \theta_W \right) \right),$$

$$F_Z^A = F_D \frac{g_A}{4} F_A. \quad (\text{A7})$$

We can generalize the formalism for any combination of isosinglet and isotriplet axial and axial-vector currents, meaning that we can get the form factor for all the possible flavor-conserving quark coupling structures of generic JJ portals.

Notice that in Table III there are no axial-vector isosinglets form factors, but they can be obtained from lattice computations.

APPENDIX B: MESON-DECAY MATRIX ELEMENTS

Here we summarize the QCD matrix elements and form factor parametrizations we use for computing decays of mesons. For the case of annihilation decays, the meson decays to DS states entirely, while for radiative decays, a heavier meson decays to lighter mesons, along with DS. The matrix element for the process can be factorized into a short distance contribution (and only involves DS matrix elements) and a long distance (QCD) contribution. For the process $\mathbf{M} \rightarrow \text{DS}$, where \mathbf{M} is the decaying meson, the full amplitude $\langle \text{DS} | \mathcal{O}_{\text{SM}} \mathcal{O}_{\text{DS}} | \mathbf{M} \rangle$ factors into $\langle 0 | \mathcal{O}_{\text{SM}} | \mathbf{M} \rangle \times \langle \text{DS} | \mathcal{O}_{\text{DS}} | 0 \rangle$ while for the process $\mathbf{H} \rightarrow \mathbf{L} + \text{DS}$, where \mathbf{H} (\mathbf{L}) are the heavy (light) SM mesons, the amplitude $\langle \mathbf{L}, \text{DS} | \mathcal{O}_{\text{SM}} \mathcal{O}_{\text{DS}} | \mathbf{H} \rangle$ factors as $\langle \mathbf{L} | \mathcal{O}_{\text{SM}} | \mathbf{H} \rangle \times \langle \text{DS} | \mathcal{O}_{\text{DS}} | 0 \rangle$. For the annihilation case, we only consider vector mesons, denoted by V .¹⁴ The SM matrix element for annihilation decay is simple, and is generically given as

$$\langle 0 | \bar{u} \gamma_\mu u | V(p) \rangle = i f_V m_V \epsilon_\mu(p), \quad (\text{B1})$$

where u are the quark spinors, $\epsilon_\mu(p)$ is the polarization vector of the vector meson V , f_V is the decay constant, and m_V is the mass of the meson.

For the radiative decay case, the SM contribution to the amplitude is the same as SM semileptonic meson decays, and is less straightforward than the annihilation case. We now give details of the matrix elements and form factors used to compute width of decays of $\mathbf{H} \rightarrow \mathbf{L} + \text{DS}$ where \mathbf{H} can be B , K and \mathbf{L} can either be a light pseudoscalar P (e.g., K , π) or vector V (e.g., K^* , ρ , ϕ).

1. Decay to pseudoscalars

For the decay of mesons to pseudoscalars P (e.g., K , π) we use the usual matrix element definitions (see Ref. [135] for example),

¹⁴For pseudoscalar mesons, the matrix element is proportional to p_μ and vanishes when contracted with the DS current, by current conservation.

$$\langle P(p_P)|V^\mu|H(p_H)\rangle = f_+(q^2)p^\mu + (f_0(q^2) - f_+(q^2))Dq^\mu, \quad (\text{B2})$$

where $V^\mu = \bar{u}_L \gamma^\mu u_H$, $p^\mu = p_H^\mu + p_P^\mu$, $D = (M_H^2 - M_P^2)/p_{\text{DS}}^2$, and $q^\mu \equiv p_{\text{DS}}^\mu = p_H^\mu - p_P^\mu$. Here, u_L (u_H) denotes a light (heavy) quark field, p_H (p_P) is the 4-momentum of the decaying heavy (light) meson with mass m_H (m_P), and $f_0(q^2), f_+(q^2)$ are dimensionless form factors which encode the strong interaction effects. For the case of $\mathcal{O}_{\text{DS}} = J_\mu^{\text{DS}}$, a conserved current, terms proportional to $q^\mu J_\mu^{\text{DS}}(q)$ vanish, so that we only need to specify the $f_+(q^2)$ form factors.

For $K^+ \rightarrow \pi^+ + \text{DS}$ decay we use the explicit form factor data points defined at each q^2 in Table IV of Ref. [136]. For $B \rightarrow \pi, K$ decays, we use the form factor definitions and values from Ref. [135],

$$f_+^{B \rightarrow \pi}(q^2) = \frac{A}{(1 - q^2/D^2)(1 - q^2/E^2)},$$

$$f_+^{B \rightarrow K}(q^2) = \frac{B}{1 - q^2/F^2} + \frac{C}{(1 - q^2/F^2)^2}, \quad (\text{B3})$$

where $A = 0.258$, $B = 0.173$, $C = 0.162$ and $D = 5.32$ GeV, $E = 6.38$ GeV, and $F = 5.41$ GeV.

2. Decay to vector mesons

For the case of B meson decaying to light vector mesons V (e.g., K^*, ρ, ϕ), the QCD matrix element is defined as (see Ref. [137], and Refs. [138–141] for more recent analyses and values of B -meson form factors)

$$\begin{aligned} \langle V(p_V)|J_\mu|B(p_B)\rangle = & -i\epsilon_{\mu\nu}^*(p_V)(m_B + m_V)A_1(q^2) \\ & + i p_\mu(\epsilon^*(p_V) \cdot q) \frac{A_2(q^2)}{m_B + m_V} \\ & + i q_\mu(\epsilon^*(p_V) \cdot q) \frac{2m_V}{q^2} \\ & \times (A_3(q^2) - A_0(q^2)) \\ & + \epsilon_{\mu\nu\rho\sigma} \epsilon^{\nu*}(p_V) p_B^\rho p_V^\sigma \frac{2V(q^2)}{m_B + m_V}, \quad (\text{B4}) \end{aligned}$$

where $J^\mu = \bar{q}_L \gamma^\mu (1 - \gamma_5) q_H$, $p^\mu = (p_B + p_V)^\mu$, and $q^\mu \equiv p_{\text{DS}}^\mu = p_B^\mu - p_V^\mu$. Here p_B (p_V) are the 4-momentum of the B (vector) meson, m_B (m_V) is the mass of the B (V) meson, ϵ^μ is the polarization vector of V , and $A_i(q^2), V(q^2)$ are the dimensionless form factors encoding strong-interaction effects. Note that, unlike the pseudoscalar case, in the case of vector meson, the $\gamma_5 \gamma^\mu$ part of the current does not vanish. Further, the third term in the above does not contribute since it is zero by current conservation, as in the pseudoscalar case.

TABLE IV. Fit parameters for the form factors defined in Eqs. (B5)–(B7) for various $B \rightarrow V$ transitions where $V = K^*, \rho, \phi$ (taken from [137]).

	r_1	r_2	m_R^2 [GeV ²]	m_{fit}^2 [GeV ²]
$V^{B \rightarrow K^*}$	0.923	-0.511	28.30	49.4
$A_1^{B \rightarrow K^*}$...	0.290	...	40.38
$A_2^{B \rightarrow K^*}$	-0.084	0.342	...	52.00
$V^{B \rightarrow \rho}$	1.045	-0.721	28.30	38.34
$A_1^{B \rightarrow \rho}$...	0.240	...	37.51
$A_2^{B \rightarrow \rho}$	0.009	0.212	...	40.82
$V^{B_s \rightarrow \phi}$	1.484	-1.049	29.38	39.52
$A_1^{B_s \rightarrow \phi}$...	0.308	...	36.54
$A_2^{B_s \rightarrow \phi}$	-0.054	0.288	...	48.94

For the form factors, we use the parametrizations as given in Ref. [137],

$$V(q^2) = \frac{r_1}{1 - q^2/m_R^2} + \frac{r_2}{1 - q^2/m_{\text{fit}}^2}, \quad (\text{B5})$$

$$A_1(q^2) = \frac{r_2}{1 - q^2/m_{\text{fit}}^2}, \quad (\text{B6})$$

$$A_2(q^2) = \frac{r_1}{1 - q^2/m_{\text{fit}}^2} + \frac{r_2}{(1 - q^2/m_{\text{fit}}^2)^2}. \quad (\text{B7})$$

We give the values of the various fit parameters for the different decays in Table IV. To evaluate further the squared matrix element, we make the calculations in the rest frame of B meson, where only the longitudinal polarization of $\epsilon_\mu^*(p_V)$ contributes. Thus, using $\epsilon_\mu^*(p_V) = (\frac{|\vec{p}_V|}{m_V}, \frac{\vec{p}_V}{|\vec{p}_V|} \frac{E_V}{m_V})$ and $p_{\text{DS}} = p_B - p_V$, we can find the production width of DS from radiative $B \rightarrow V$ decays.

3. Sensitivity to nonconformal contributions in K -, B - meson decays

For the case of DS production from irrelevant portals, we expect the production cross section to grow with p_{DS}^2 which is necessary to make the contribution away from the IR threshold more important, and is needed for usefulness of our model-agnostic approach. For radiative meson decays however, the DS production cross section does not grow as p_{DS}^2 , but is rather flat, up to the kinematic threshold. For annihilation decays, p_{DS}^2 is fixed and equals the parent meson mass-squared. To justify our model-agnostic approach we need to ensure we are away from the thresholds, which must be imposed as a self-consistent criteria.

The overall signal is obtained by an integral over the range of allowed p_{DS} . The kinematic condition $p_{\text{DS}}/(n_{\text{LDSP}} \Lambda_{\text{IR}}) \geq 1$ is always stronger than the condition to be in the conformal regime $p_{\text{DS}}/\Lambda_{\text{IR}} \gtrsim 1$. We also need to make sure that the

relevant Λ_{IR} probed in the experiment under consideration (which depends on the portal and the lifetime) is away from the kinematic threshold $M - m$. For Λ_{IR} close to $M - m$, a small change in the lower limit of p_{DS} integration would have a bigger impact, but this does not happen in the cases we consider. For $B \rightarrow K + \text{DS}$, $(m_B - m_K) \sim 5$ GeV is larger than typical Λ_{IR} probed which is 0.001–1 GeV, while for $K \rightarrow \pi + \text{DS}$, $(m_K - m_\pi) \sim 0.4$ GeV, while the typical Λ_{IR} probed is 0.001–0.1 GeV.

APPENDIX C: PROBABILITY OF DECAY

To compute the number of signal events, we need to calculate the number of LDSPs that decay inside the detector. To correctly compute this quantity, the differential cross section (in both energy and angle) of LDSP production must be convoluted with the probability P for at least an LDSP to decay inside the detector. The probability P_1 for a particle to decay inside the detector can be roughly estimated as the probability $P_{1,\text{dec}}$ to decay within the radial distance at which the detector is, multiplied with ϵ_{geo} , the geometric acceptance accounting for the particles flying in the detector direction,

$$P_1 \approx \epsilon_{\text{geo}} P_{1,\text{dec}}. \quad (\text{C1})$$

In a simplified setup (see Fig. 6), $P_{1,\text{dec}}$ can be estimated to be

$$P_{1,\text{dec}} = \exp\left(-\frac{l}{c\tau(\gamma\beta)_{\text{LDSP}}}\right) - \exp\left(-\frac{l+d}{c\tau(\gamma\beta)_{\text{LDSP}}}\right), \quad (\text{C2})$$

where l and d are the distance of the detector from the target location and the length of the detector respectively,¹⁵ γ_{LDSP} is the (energy-dependent) Lorentz factor of the decaying particle, β_{LDSP} is its velocity and τ is its proper lifetime. In general, for more than 1 LDSP, a slightly more refined procedure accounting for the presence of multiple particles that can go inside the detector is needed. We now discuss our procedure for multiple-particle events. We consider both the weakly coupled case, where $n_{\text{LDSP}} = 2$, and a generic strongly coupled case in which $n_{\text{LDSP}} \sim \mathcal{O}(10)$ are produced.

For the weakly coupled case, the directions of the two LDSPs are fully correlated, at fixed p_{DS} , by momentum conservation. If the direction of the DS system (i.e., \vec{p}_{DS}) does not intersect the detector, at most one of the two LDSPs can have the correct direction to hit the detector.

¹⁵Notice that the distance should be a function of the direction of the LDSP for generic geometry of the detector. We will work under the assumption of spherical detectors, where the distances appearing in P are independent of the line of flight. This is a good approximation for neutrino experiments, since they consist of small boxes [of size $\mathcal{O}(10)$ m] positioned far [$\mathcal{O}(10^3)$ m] from the target. The error due to the geometrical approximation is therefore of order 1% (the ratio of the typical distances).

Indeed, the two particles have opposite azimuthal angle ϕ (computed with respect to the DS momentum in the lab frame), so that only one particle at most can travel to the correct side. Therefore, the probability to have at least one particle decaying inside the detector is [in the notation of Eq. (9)],

$$\epsilon_{\text{geo}} P_{\text{decay}} = 2\epsilon_{\text{geo}} P_{1,\text{dec}},$$

where the 2 reflects the fact that ϵ_{geo} has been computed for a single particle only, and there are two possible particles that can decay in the detector (or in other words, we only considered one of the particles to be aimed toward the detector, but the opposite direction is also a valid choice given the presence of the other LDSP).

If instead the DS direction intersects the detector there is the possibility for both LDSPs to fall inside the detector. This can happen only if the LDSP velocity in the DS frame is slower than the velocity of the DS system in the lab frame, so that the LDSP traveling in the direction $\theta > \pi/2$ (the backward direction) in the DS frame is boosted in the forward direction. The events in which the backward LDSP gets boosted forward in the lab frame are only significant if the DS system velocity is much larger than the velocity of the LDSP in the DS frame. Given that typically $n_{\text{LDSP}}\Lambda_{\text{IR}} \ll \sqrt{p_{\text{DS}}^2}$, and that the bulk of the cross section is dominated by high p_{DS}^2 events, the events will have fast LDSPs in the DS frame. Therefore, events with two particles boosted forward are negligible, and the formula to have a particle decaying inside the detector holds unchanged.

In the strongly coupled case, where n_{LDSP} is large, we can neglect the fact that the momenta are correlated by momentum conservation. Therefore, we will assume that the n_{LDSP} particles share democratically the energy in the rest frame, and that the directions are independent samples from an isotropic distribution in such a frame. In this scenario, multiple particles can in principle have the correct direction to get inside the detector in the lab frame. Using the probability for a single particle to decay inside the detector from Eq. (C2), and recalling that this probability is very small (both $P_{1,\text{dec}}$ and typical ϵ_{geo} are small), it follows that events in which multiple particles simultaneously decay inside the detector are very rare. Therefore, we can approximate the probability for a single event to contain at least one particle decaying inside the detector as

$$\epsilon_{\text{geo}} P_{\text{decay}} = n_{\text{LDSP}} \epsilon_{\text{geo}} P_{1,\text{dec}}.$$

Notice that in the end (although for slightly different reason) this formula matches the one in the weakly coupled case for $n_{\text{LDSP}} = 2$.

APPENDIX D: ϵ_{geo} ESTIMATES

In order to compute the signal events, we have to compute the fraction of events that contain at least one LDSP with the direction intersecting the detector. Given the assumptions outlined in Appendix C, we define ϵ_{geo} as the probability for a single particle to have such a direction. We work under the assumption that LDSPs are produced isotropically in the DS rest frame, the frame in which only the time component of p_{DS} is nonzero. We will also assume that in this frame, all the LDSPs produced share the same energy $p_{\text{DS}}^0/n_{\text{LDSP}}$. We take the beam direction to be aligned to the z -axis ($\theta = 0$ in polar coordinates), and ϕ as the azimuthal angle (measured in the plane orthogonal to the beam line).

In order to get an estimate for the angular coverage of the neutrino detector, we work under the simplifying assumption that the detector surface lies on a 2D plane orthogonal to the beam (for the off-axis case, the angle is very small) and all points on the detector are at the same distance from the interaction point. This is a good approximation for neutrino detectors since corrections are of order $(d/l)^2 \ll 1$, where l is the distance at which the detector is placed, and d is the typical size of the detector. There are two relevant cases where this approximation fails.

For a closer experiment like SHiP, l and d are of same order, and the proposed shape is a rectangle of dimensions $(a, 2a)$, $a = 5$ m. We will compute the angular acceptance by taking the largest side of the rectangle and requiring the trajectory to intersect the first layer, which gives $\theta_{\text{acc}} = \text{ArcTan}(2a/2l)$. Being more specific about the shape of the detector does not change the estimate significantly since the DS is produced very boosted, with a small opening angle, and is fully covered by the extent of the detector. The geometric acceptances are close to unity, see Table I.

For MicroBooNE, the produced kaons are at rest in the lab frame, and not collimated, so that the DS is produced isotropic in the lab frame. In this case, the above approximation again does not hold. We calculate ϵ_{geo} without this approximation for MicroBooNE.

To calculate ϵ_{geo} , we need to find the overlap of the LDSPs with the detector, which is easier to do in the $*$ frame. By the assumption of isotropic decay, the LDSP directions are distributed uniformly in the $(\cos\theta^*, \phi^*)$ plane. In this plane ϵ_{geo} is the area (normalized by $1/4\pi$) that corresponds to lab frame configurations in which the LDSP falls inside the detector. This in general depends on p_{DS} and n_{LDSP} through the boost factor.

To be precise, we need to introduce some notation. There are two frames to consider, the lab frame, and the DS frame (the one where p_{DS} only has a time component). We refer to the DS frame as the $*$ frame, in what follows. The angles and the boost factors with a $*$ superscript are defined in the DS frame, and without a $*$ superscript are in the lab frame.

The detector direction, the DS direction, and the LDSP direction are all defined by a θ , and the last two depend in general on kinematical quantities such as p_{DS} and other process-dependent quantities e.g., z in bremsstrahlung. We take the detector to be positioned at $\theta = \theta_{\text{det}}$. The accepted directions form a cone around it with an angle θ_{acc} around it. The DS system may be along the beam line or at an angle from it. We define the DS system to be at $\theta = \theta_{\text{DS}}$. Finally, the LDSP is at an angle θ_{LDSP} with respect to the DS system, so that the DS direction has a cone of angle θ_{LDSP} around it. It is useful to also define the DS direction with respect to the detector, which is denoted by $\theta_{\text{det}}^{\text{eff}} \approx (\theta_{\text{DS}}^2 + \theta_{\text{det}}^2 - 2\theta_{\text{det}}\theta_{\text{DS}}\cos\phi_{\text{DS}})^{1/2}$. All these quantities are shown in Fig. 9 (drawn not to scale), and are defined in the lab frame.

We additionally denote the LDSP in the $*$ frame to be at (θ^*, ϕ^*) . For a given p_{DS} and θ^* , θ_{LDSP} is given as

$$\tan\theta_{\text{LDSP}} = \frac{\sin\theta^*}{\gamma_{\text{DS}}(\beta_{\text{DS}}/\beta_{\text{LDSP}}^* + \cos\theta^*)}, \quad (\text{D1})$$

where β_{DS} , γ_{DS} are the boost factors to go from $*$ to the lab frame and β_{LDSP}^* is the velocity of the LDSP with respect to the $*$ frame. Note that the boosts along the DS direction do not change ϕ . Given the width θ_{LDSP} of the lab cone there is only a range of $\phi_* = \phi_{\text{lab}}$ for which the LDSP direction intersects the detector. ϵ_{geo} is then the area (computed in the $*$ frame) of the overlap between the circular detector and all the possible LDSP cones in the lab frame, in general a function of z and of p_{DS} . To compute such area, we will make a linearization approximation; we will consider that in the $*$ frame the shape of allowed $\cos\theta^* - \phi^*$ is bound by straight lines and not curved ones. This approximation is expected to hold at the 10% level for the relevant boosts.

We now discuss the details specific to the three production modes considered in this work.

1. ϵ_{geo} for meson-decay production

In order to find ϵ_{geo} for LDSPs in the meson case, we use the prescription explained in the section before. For this, we must boost the DS kinematic variables twice; once from the meson rest frame to the lab frame, then from the lab frame to the DS rest frame (or $*$ frame). We will denote all quantities in the meson rest frame with a subscript 0.

The energy and 3-momentum of the DS can be easily computed using momentum conservation in the parent meson rest frame as shown in Eq. (19). Using the parent meson boost γ_M and velocity v_M (obtained from simulation), we can find the DS lab frame variables and express them as a function of $\cos\theta_0$ (angle of DS with meson flight direction in the meson rest frame, DS is isotropic in this) and p_{DS}^2 i.e., $E_{\text{DS}}^{\text{lab}}(\cos\theta_0, p_{\text{DS}}^2)$, $|\vec{p}_{\text{DS}}^{\text{lab}}|(\cos\theta_0, p_{\text{DS}}^2)$.

In order to compute ϵ_{geo} we must find the θ^* values corresponding to the angular coverage of the detector. For

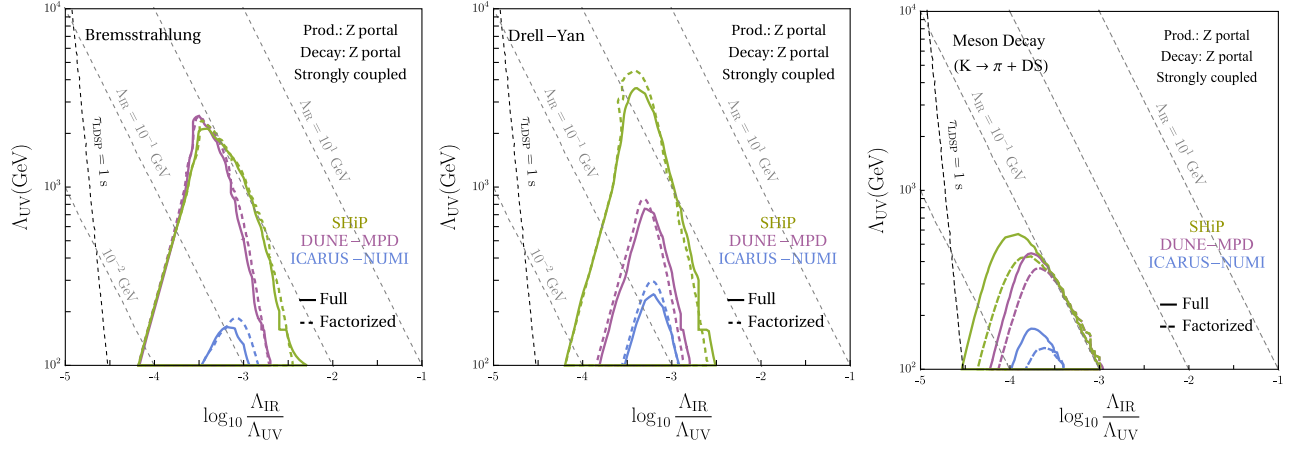


FIG. 10. Comparison of the factorized (dashed) vs full (solid) approach for production and decay through Z portal of a strongly interacting DS, for various production modes: DB (left), DY (mid), and MD (right). The experiments considered are ICARUS-NuMI (blue), DUNE-MPD (purple), and SHiP (green).

and $\theta_{DS} = 0$ by momentum conservation, which simplifies the equation above.

2. ϵ_{geo} for Drell-Yan

In DY mode the DS system is by construction directed along the original beam line (the z direction). The boost of the DS system is $\gamma_{DS} = E_{DS}/\sqrt{p_{DS}^2}$, where E_{DS} is the energy of the DS system in the lab frame. While p_{DS} is already one of the variables used in DY production, E_{DS} is computed by first getting the DS energy in the DY CM frame, $E_{DS,CM} = \sqrt{s}(x + p_{DS}^2/(sx))/2$ using the DY variables introduced in Sec. III B, and then boosting it to the lab frame with Lorentz parameter $\gamma_{CM} = E_{\text{beam}}/\sqrt{s}$. This gives $E_{DS} = xE_{\text{beam}}$. Putting everything together, the boost factor to go from the lab frame to the DS system is $\gamma_{DS} = xE_{\text{beam}}/\sqrt{p_{DS}^2}$. To get the average value of ϵ_{geo} , we compute the integral as in Eq. (D2) with the appropriate distribution where now we average over the kinematic variables x and p_{DS}^2 .

3. ϵ_{geo} for dark bremsstrahlung

DB mode is different from the DY mode because in general the DS system will be produced at an angle with the beam-line. We will call θ_{DS}, ϕ_{DS} the pair of angles indicating the direction of the radiated DS system relative to the beam line. Using the kinematic variables introduced in Sec. III C, we have $\theta_{DS} = \tan^{-1}(p_T/zE_{\text{beam}})$. The boost factor to go from the lab to the DS frame is $\gamma_{DS} \approx zE_{\text{beam}}/\sqrt{p_{DS}^2}$. As before, to get the average value of ϵ_{geo} , we average over the kinematic quantities z, p_T^2 , and p_{DS}^2 . We have checked that these values are in good agreement with the ϵ_{geo} for the average DS angle, as expected for very collimated DS excitations.

APPENDIX E: FACTORIZATION APPROXIMATION

In this appendix we compare the approximate method we delineated in the main text with the correct procedure of doing the integral of the product of the differential quantities (ϵ_{geo} , cross section, decay probability) without factorizing them.

In order to get the correct number of signal events, the differential cross section $d\sigma/d\lambda_i$ must be folded with $n_{\text{LDSP}}\epsilon_{\text{geo}}P_{1,\text{dec}}$, all of which are a function of kinematic variables λ_i ,

$$N_S(\Lambda_{\text{IR}}, \Lambda_{\text{UV}}) = \frac{N_{\text{POT}}}{\sigma_{pN}} \int d\lambda_i \frac{d\sigma}{d\lambda_i} n_{\text{LDSP}}\epsilon_{\text{geo}}P_{1,\text{dec}}. \quad (\text{E1})$$

We call this procedure the full approach. We compute the integrals numerically, using the CUBA integration tools [142]. The method used in the main text is done instead by replacing the full integral with the average ϵ_{geo} and using average kinematic quantities to estimate the decay probability and the average n_{LDSP} ,¹⁶

$$N_S \approx N_{\text{POT}} \frac{\sigma_S}{\sigma_{pN}} n_{\text{LDSP}}(\langle\lambda_i\rangle) P_{1,\text{dec}}(\langle\lambda_i\rangle) \langle\epsilon_{\text{geo}}\rangle, \quad (\text{E2})$$

where the average is defined by a weighted integral over the differential cross section/decay width [e.g. see Eq. (D2)]. We call this approach the factorized approach. The advantage of this approach is that the production integral must be done only once, and not repeated for each $(\Lambda_{\text{IR}}, \Lambda_{\text{UV}})$ pair; the dependence on them is essentially factorized. In particular, for a fixed production portal, this allows changing the decay portal, without having to redo the production integral from scratch. This is particularly useful when

¹⁶We have checked that this procedure agrees very well (percent level) to a true average of n_{LDSP} .

exploring all the various combinations of production and decay portals.

To compare the factorized and the full approach, and to show that the factorized approach is very efficient, in Fig. 10 we show the comparison for DB, DY, and radiative meson decay $K \rightarrow \pi + DS$, for the strongly coupled DS, for

DUNE, SHiP, and ICARUS. The two approaches are in very good agreement. The factorized approach is conservative at most, and can miss rare events appearing in the tails of distributions (see e.g., Ref. [143]). However for the purposes of the present work, the factorized approach suffices.

-
- [1] M. J. Strassler and K. M. Zurek, Echoes of a hidden valley at hadron colliders, *Phys. Lett. B* **651**, 374 (2007).
- [2] Y. Kahn, G. Krnjaic, N. Tran, and A. Whitbeck, M3: A new muon missing momentum experiment to probe muon $g-2$ and dark matter at Fermilab, *J. High Energy Phys.* **09** (2018) 153.
- [3] H. Sieber, D. Banerjee, P. Crivelli, E. Depero, S. N. Gninenko, D. V. Kirpichnikov, M. M. Kirsanov, V. Poliakov, and L. Molina Bueno, Prospects in the search for a new light Z' boson with the NA64 μ experiment at the CERN SPS, *Phys. Rev. D* **105**, 052006 (2022).
- [4] C. Cesarotti, S. Homiller, R. K. Mishra, and M. Reece, Probing New Gauge Forces with a High-Energy Muon Beam Dump, *Phys. Rev. Lett.* **130**, 071803 (2023).
- [5] B. Batell, N. Blinov, C. Hearty, and R. McGehee, Exploring dark sector portals with high intensity experiments, in *2022 Snowmass Summer Study* (2022) [arXiv:2207.06905].
- [6] B. Batell *et al.*, Dark sector studies with neutrino beams, in *2022 Snowmass Summer Study* (2022) [arXiv:2207.06898].
- [7] G. Krnjaic *et al.*, A snowmass whitepaper: Dark matter production at intensity-frontier experiments, arXiv:2207.00597.
- [8] J. Berger *et al.*, Snowmass 2021 white paper: Cosmogenic dark matter and exotic particle searches in neutrino experiments, in *2022 Snowmass Summer Study* (2022) [arXiv:2207.02882].
- [9] S. Gori *et al.*, Dark sector physics at high-intensity experiments, arXiv:2209.04671.
- [10] P. Coloma, L. W. Koerner, I. M. Shoemaker, and J. Yu, Neutrino frontier topical group report (NF03): Physics beyond the standard model, arXiv:2209.10362.
- [11] M. Antonello *et al.* (MicroBooNE, LAr1-ND, and ICARUS-WA104 Collaborations), A proposal for a three detector short-baseline neutrino oscillation program in the Fermilab booster neutrino beam, arXiv:1503.01520.
- [12] R. Acciarri *et al.* (MicroBooNE Collaboration), Design and construction of the MicroBooNE detector, *J. Instrum.* **12**, P02017 (2017).
- [13] P. A. Machado, O. Palamara, and D. W. Schmitz, The short-baseline neutrino program at Fermilab, *Annu. Rev. Nucl. Part. Sci.* **69**, 363 (2019).
- [14] J. M. Berryman, A. de Gouvea, P. J. Fox, B. J. Kayser, K. J. Kelly, and J. L. Raaf, Searches for decays of new particles in the DUNE multi-purpose near detector, *J. High Energy Phys.* **02** (2020) 174.
- [15] G. Chauhan, P. S. B. Dev, and X.-J. Xu, Probing the ν_R -philic Z' at DUNE near detectors, *Phys. Lett. B* **841**, 137907 (2023).
- [16] B. Batell, J. Berger, and A. Ismail, Probing the Higgs portal at the Fermilab short-baseline neutrino experiments, *Phys. Rev. D* **100**, 115039 (2019).
- [17] P. Ballett, M. Hostert, S. Pascoli, Y. F. Perez-Gonzalez, Z. Tabrizi, and R. Zukanovich Funchal, Neutrino trident scattering at near detectors, *J. High Energy Phys.* **01** (2019) 119.
- [18] D. Curtin *et al.*, Long-lived particles at the energy frontier: The MATHUSLA physics case, *Rep. Prog. Phys.* **82**, 116201 (2019).
- [19] J. Beacham *et al.*, Physics beyond colliders at CERN: Beyond the standard model working group report, *J. Phys. G* **47**, 010501 (2020).
- [20] R. Essig *et al.*, Working group report: New light weakly coupled particles, in *Community Summer Study 2013: Snowmass on the Mississippi* (2013) [arXiv:1311.0029].
- [21] J. Alexander *et al.*, Dark sectors 2016 workshop: Community report, arXiv:1608.08632.
- [22] B. Döbrich, J. Jaeckel, F. Kahlhoefer, A. Ringwald, and K. Schmidt-Hoberg, ALPtraum: ALP production in proton beam dump experiments, *J. High Energy Phys.* **02** (2016) 018.
- [23] B. Döbrich, J. Jaeckel, and T. Spadaro, Light in the beam dump—ALP production from decay photons in proton beam-dumps, *J. High Energy Phys.* **05** (2019) 213; *J. High Energy Phys.* **10** (2020) 46.
- [24] W. Altmannshofer, S. Gori, and D. J. Robinson, Constraining axionlike particles from rare pion decays, *Phys. Rev. D* **101**, 075002 (2020).
- [25] K. Asai, S. Iwamoto, Y. Sakaki, and D. Ueda, New physics searches at the ILC positron and electron beam dumps, *J. High Energy Phys.* **09** (2021) 183.
- [26] E. Bertuzzo, A. L. Foguel, G. M. Salla, and R. Z. Funchal, New Limits on Leptophilic ALPs and Majorons from ArgoNeuT, *Phys. Rev. Lett.* **130**, 171801 (2023).
- [27] L. Harland-Lang, J. Jaeckel, and M. Spannowsky, A fresh look at ALP searches in fixed target experiments, *Phys. Lett. B* **793**, 281 (2019).
- [28] P. deNiverville, H.-S. Lee, and M.-S. Seo, Implications of the dark axion portal for the muon $g-2$, B factories, fixed target neutrino experiments, and beam dumps, *Phys. Rev. D* **98**, 115011 (2018).
- [29] R. T. Co, S. Kumar, and Z. Liu, Enhancing searches for heavy QCD axions via dimuon final states, *J. High Energy Phys.* **02** (2023) 111.

- [30] R. Contino, K. Max, and R. K. Mishra, Searching for elusive dark sectors with terrestrial and celestial observations, *J. High Energy Phys.* **06** (2021) 127.
- [31] C. Arina, J. Hajer, and P. Klose, Portal effective theories: A framework for the model independent description of light hidden sector interactions, *J. High Energy Phys.* **09** (2021) 063.
- [32] H.-C. Cheng, L. Li, and E. Salvioni, A theory of dark pions, *J. High Energy Phys.* **01** (2022) 122.
- [33] L. Darmé, S. A. R. Ellis, and T. You, Light dark sectors through the fermion portal, *J. High Energy Phys.* **07** (2020) 053.
- [34] E. Bertuzzo and M. Taoso, Probing light dark scalars with future experiments, *J. High Energy Phys.* **03** (2021) 272.
- [35] A. Berlin, P. deNiverville, A. Ritz, P. Schuster, and N. Toro, Sub-GeV dark matter production at fixed-target experiments, *Phys. Rev. D* **102**, 095011 (2020).
- [36] A. Carmona, C. Scherb, and P. Schwaller, Charming ALPs, *J. High Energy Phys.* **08** (2021) 121.
- [37] D. Barducci, E. Bertuzzo, M. Taoso, and C. Toni, Probing right-handed neutrinos dipole operators, *J. High Energy Phys.* **03** (2023) 239.
- [38] E. Bernreuther, K. Böse, T. Ferber, C. Hearty, F. Kahlhoefer, A. Morandini, and K. Schmidt-Hoberg, Forecasting dark showers at Belle II, *J. High Energy Phys.* **12** (2022) 005.
- [39] P. Schwaller, D. Stolarski, and A. Weiler, Emerging jets, *J. High Energy Phys.* **05** (2015) 059.
- [40] D. E. Kaplan, M. A. Luty, and K. M. Zurek, Asymmetric dark matter, *Phys. Rev. D* **79**, 115016 (2009).
- [41] K. M. Zurek, Asymmetric dark matter: Theories, signatures, and constraints, *Phys. Rep.* **537**, 91 (2014).
- [42] K. Petraki and R. R. Volkas, Review of asymmetric dark matter, *Int. J. Mod. Phys. A* **28**, 1330028 (2013).
- [43] R. Foot, H. Lew, and R. R. Volkas, A model with fundamental improper space-time symmetries, *Phys. Lett. B* **272**, 67 (1991).
- [44] R. Foot, Mirror dark matter: Cosmology, galaxy structure and direct detection, *Int. J. Mod. Phys. A* **29**, 1430013 (2014).
- [45] Z. Chacko, H.-S. Goh, and R. Harnik, The Twin Higgs: Natural Electroweak Breaking from Mirror Symmetry, *Phys. Rev. Lett.* **96**, 231802 (2006).
- [46] B. Batell, J. Berger, L. Darmé, and C. Frugiuele, Inelastic dark matter at the Fermilab short baseline neutrino program, *Phys. Rev. D* **104**, 075026 (2021).
- [47] P. deNiverville and C. Frugiuele, Hunting sub-GeV dark matter with the NO ν A near detector, *Phys. Rev. D* **99**, 051701 (2019).
- [48] L. Buonocore, C. Frugiuele, and P. deNiverville, Hunt for sub-GeV dark matter at neutrino facilities: A survey of past and present experiments, *Phys. Rev. D* **102**, 035006 (2020).
- [49] P. deNiverville, M. Pospelov, and A. Ritz, Observing a light dark matter beam with neutrino experiments, *Phys. Rev. D* **84**, 075020 (2011).
- [50] B. Batell, P. deNiverville, D. McKeen, M. Pospelov, and A. Ritz, Leptophobic dark matter at neutrino factories, *Phys. Rev. D* **90**, 115014 (2014).
- [51] B. Batell, M. Pospelov, and A. Ritz, Exploring portals to a hidden sector through fixed targets, *Phys. Rev. D* **80**, 095024 (2009).
- [52] P. deNiverville, C.-Y. Chen, M. Pospelov, and A. Ritz, Light dark matter in neutrino beams: Production modelling and scattering signatures at MiniBooNE, T2K and SHiP, *Phys. Rev. D* **95**, 035006 (2017).
- [53] P. deNiverville, D. McKeen, and A. Ritz, Signatures of sub-GeV dark matter beams at neutrino experiments, *Phys. Rev. D* **86**, 035022 (2012).
- [54] D. Tucker-Smith and N. Weiner, Inelastic dark matter, *Phys. Rev. D* **64**, 043502 (2001).
- [55] M. W. Winkler, Decay and detection of a light scalar boson mixing with the Higgs boson, *Phys. Rev. D* **99**, 015018 (2019).
- [56] P. Ilten, Y. Soreq, M. Williams, and W. Xue, Serendipity in dark photon searches, *J. High Energy Phys.* **06** (2018) 004.
- [57] C. Baruch, P. Ilten, Y. Soreq, and M. Williams, Axial vectors in DarkCast, *J. High Energy Phys.* **11** (2022) 124.
- [58] B. R. Webber, Average multiplicities in jets, *Phys. Lett.* **143B**, 501 (1984).
- [59] C. Cesarotti, M. Reece, and M. J. Strassler, Spheres to jets: Tuning event shapes with 5d simplified models, *J. High Energy Phys.* **05** (2021) 096.
- [60] N. Blinov, E. Kowalczyk, and M. Wynne, Axion-like particle searches at DarkQuest, *J. High Energy Phys.* **02** (2022) 036.
- [61] A. J. Buras, Weak Hamiltonian, CP violation and rare decays, in *Les Houches Summer School in Theoretical Physics, Session 68: Probing the Standard Model of Particle Interactions* (1998), pp. 281–539 [arXiv:hep-ph/9806471].
- [62] T. Inami and C. S. Lim, Effects of superheavy quarks and leptons in low-energy weak processes $K_L \rightarrow \mu\bar{\mu}$, $K^+ \rightarrow \pi^+ \nu\bar{\nu}$ and $K^0 \leftrightarrow \bar{K}^0$, *Prog. Theor. Phys.* **65**, 297 (1981); *Prog. Theor. Phys.* **65**, 1772(E) (1981).
- [63] C. Ahdida *et al.* (SHiP Collaboration), Sensitivity of the SHiP experiment to heavy neutral leptons, *J. High Energy Phys.* **04** (2019) 077.
- [64] J. L. Feng, I. Galon, F. Kling, and S. Trojanowski, Dark Higgs bosons at the forward search experiment, *Phys. Rev. D* **97**, 055034 (2018).
- [65] D. Gorbunov, I. Krasnov, Y. Kudenko, and S. Suvorov, Heavy neutral leptons from kaon decays in the SHiP experiment, *Phys. Lett. B* **810**, 135817 (2020).
- [66] I. Boiarska, K. Bondarenko, A. Boyarsky, V. Gorkavenko, M. Ovchinnikov, and A. Sokolenko, Phenomenology of GeV-scale scalar portal, *J. High Energy Phys.* **11** (2019) 162.
- [67] W. Altmannshofer, J. A. Dror, and S. Gori, New Insights Into Axion-Lepton Interactions, *Phys. Rev. Lett.* **130**, 241801 (2023).
- [68] V. Barger, C.-W. Chiang, W.-Y. Keung, and D. Marfatia, Constraint on Parity-Violating Muonic Forces, *Phys. Rev. Lett.* **108**, 081802 (2012).
- [69] M. Hostert, K. Kaneta, and M. Pospelov, Pair production of dark particles in meson decays, *Phys. Rev. D* **102**, 055016 (2020).

- [70] A. Aduszkiewicz *et al.* (NA61/SHINE Collaboration), Measurement of ϕ meson production in $p + p$ interactions at 40, 80 and 158 GeV/c with the NA61/SHINE spectrometer at the CERN SPS, *Eur. Phys. J. C* **80**, 199 (2020).
- [71] K. Kovarik *et al.*, nCTEQ15—Global analysis of nuclear parton distributions with uncertainties in the CTEQ framework, *Phys. Rev. D* **93**, 085037 (2016).
- [72] D. B. Clark, E. Godat, and F. I. Olness, ManeParse: A *mathematica* reader for parton distribution functions, *Comput. Phys. Commun.* **216**, 126 (2017).
- [73] S. Foroughi-Abari and A. Ritz, Dark sector production via proton bremsstrahlung, *Phys. Rev. D* **105**, 095045 (2022).
- [74] M. Shifman, A. Vainshtein, and V. Zakharov, Remarks on Higgs-boson interactions with nucleons, *Phys. Lett.* **78B**, 443 (1978).
- [75] T. Feuster and U. Mosel, Photon- and meson-induced reactions on the nucleon, *Phys. Rev. C* **59**, 460 (1999).
- [76] A. S. Carroll *et al.*, Absorption cross-sections of π^\pm , K^\pm , p and \bar{p} on nuclei between 60 GeV/c and 280 GeV/c, *Phys. Lett.* **80B**, 319 (1979).
- [77] P. Abratenko *et al.* (MicroBooNE Collaboration), Search for a Higgs Portal Scalar Decaying to Electron-Positron Pairs in the MicroBooNE Detector, *Phys. Rev. Lett.* **127**, 151803 (2021).
- [78] F. Bergsma *et al.* (CHARM Collaboration), Search for axion like particle production in 400-GeV proton—copper interactions, *Phys. Lett.* **157B**, 458 (1985).
- [79] J. Dorenbosch *et al.* (CHARM Collaboration), A search for decays of heavy neutrinos in the mass range 0.5–2.8 GeV, *Phys. Lett.* **166B**, 473 (1986).
- [80] S. N. Gninenko, Constraints on sub-GeV hidden sector gauge bosons from a search for heavy neutrino decays, *Phys. Lett. B* **713**, 244 (2012).
- [81] J. Bian, Measurement of neutrino-electron elastic scattering at NOvA near detector, in *Meeting of the APS Division of Particles and Fields* (2017) [arXiv:1710.03428].
- [82] B. Abi *et al.* (DUNE Collaboration), Deep underground neutrino experiment (DUNE), far detector technical design report, volume II: DUNE physics, arXiv:2002.03005.
- [83] M. Anelli *et al.* (SHiP Collaboration), A facility to Search for Hidden Particles (SHiP) at the CERN SPS, arXiv:1504.04956.
- [84] G. Marocco and S. Sarkar, Blast from the past: Constraints on the dark sector from the BEBC WA66 beam dump experiment, *SciPost Phys.* **10**, 043 (2021).
- [85] R. Barouki, G. Marocco, and S. Sarkar, Blast from the past II: Constraints on heavy neutral leptons from the BEBC WA66 beam dump experiment, *SciPost Phys.* **13**, 118 (2022).
- [86] A. A. Aguilar-Arevalo *et al.* (MiniBooNE DM Collaboration), Dark matter search in nucleon, pion, and electron channels from a proton beam dump with MiniBooNE, *Phys. Rev. D* **98**, 112004 (2018).
- [87] P. P. Filip (NOvA Collaboration), Hunting for light dark matter with the NOvA detector, in *Proceedings of the International Conference on Neutrinos and Dark Matter* (2020).
- [88] A. Ariga *et al.* (FASER Collaboration), FASER’s physics reach for long-lived particles, *Phys. Rev. D* **99**, 095011 (2019).
- [89] H. Abreu *et al.* (FASER Collaboration), Detecting and studying high-energy collider neutrinos with FASER at the LHC, *Eur. Phys. J. C* **80**, 61 (2020).
- [90] H. Abreu *et al.* (FASER Collaboration), First Direct Observation of Collider Neutrinos with FASER at the LHC, *Phys. Rev. Lett.* **131**, 031801 (2023).
- [91] G. Acampora *et al.* (SND@LHC Collaboration), SND@LHC: The scattering and neutrino detector at the LHC, arXiv:2210.02784.
- [92] L. A. Anchordoqui *et al.*, The forward physics facility: Sites, experiments, and physics potential, *Phys. Rep.* **968**, 1 (2022).
- [93] J. L. Feng, I. Galon, F. Kling, and S. Trojanowski, Forward search experiment at the LHC, *Phys. Rev. D* **97**, 035001 (2018).
- [94] F. Kling and S. Trojanowski, Heavy neutral leptons at FASER, *Phys. Rev. D* **97**, 095016 (2018).
- [95] A. Berlin and F. Kling, Inelastic dark matter at the LHC lifetime frontier: ATLAS, CMS, LHCb, CODEX-b, FASER, and MATHUSLA, *Phys. Rev. D* **99**, 015021 (2019).
- [96] J. L. Feng, I. Galon, F. Kling, and S. Trojanowski, Axion-like particles at FASER: The LHC as a photon beam dump, *Phys. Rev. D* **98**, 055021 (2018).
- [97] A. Boyarsky, O. Mikulenko, M. Ovchinnikov, and L. Shchutka, Searches for new physics at SND@LHC, *J. High Energy Phys.* **03** (2022) 006.
- [98] T. Kuwahara and S.-R. Yuan, Dark vector mesons at LHC forward detector searches, *J. High Energy Phys.* **06** (2023) 208.
- [99] B. Batell, J. L. Feng, M. Fieg, A. Ismail, F. Kling, R. M. Abraham, and S. Trojanowski, Hadrophilic dark sectors at the Forward Physics Facility, *Phys. Rev. D* **105**, 075001 (2022).
- [100] B. Batell, J. L. Feng, and S. Trojanowski, Detecting dark matter with far-forward emulsion and liquid argon detectors at the LHC, *Phys. Rev. D* **103**, 075023 (2021).
- [101] V. Brdar, B. Dutta, W. Jang, D. Kim, I. M. Shoemaker, Z. Tabrizi, A. Thompson, and J. Yu, Axionlike Particles at Future Neutrino Experiments: Closing the Cosmological Triangle, *Phys. Rev. Lett.* **126**, 201801 (2021).
- [102] K. J. Kelly, S. Kumar, and Z. Liu, Heavy axion opportunities at the DUNE near detector, *Phys. Rev. D* **103**, 095002 (2021).
- [103] R. Essig, R. Harnik, J. Kaplan, and N. Toro, Discovering new light states at neutrino experiments, *Phys. Rev. D* **82**, 113008 (2010).
- [104] S. Foroughi-Abari and A. Ritz, LSND constraints on the Higgs portal, *Phys. Rev. D* **102**, 035015 (2020).
- [105] C. Athanassopoulos *et al.* (LSND Collaboration), Results on $\nu_\mu \rightarrow \nu_e$ oscillations from pion decay in flight neutrinos, *Phys. Rev. C* **58**, 2489 (1998).
- [106] A. Aguilar-Arevalo *et al.* (LSND Collaboration), Evidence for neutrino oscillations from the observation of $\bar{\nu}_e$ appearance in a $\bar{\nu}_\mu$ beam, *Phys. Rev. D* **64**, 112007 (2001).

- [107] M. A. Acero *et al.* (NOvA Collaboration), Measurement of the ν_e -Nucleus Charged-Current Double-Differential Cross Section at $\langle E_\nu \rangle = 2.4$ GeV using NOvA, *Phys. Rev. Lett.* **130**, 051802 (2023).
- [108] A. Bhattarai, V. Brdar, B. Dutta, W. Jang, D. Kim, I. M. Shoemaker, Z. Tabrizi, A. Thompson, and J. Yu, BSM targets at a target-less DUNE, *Phys. Rev. D* **107**, 055043 (2023).
- [109] W. Altmannshofer, S. Gori, J. Martín-Albo, A. Sousa, and M. Wallbank, Neutrino tridents at DUNE, *Phys. Rev. D* **100**, 115029 (2019).
- [110] Y. Gershtein, S. Knapen, and D. Redigolo, Probing naturally light singlets with a displaced vertex trigger, *Phys. Lett. B* **823**, 136758 (2021).
- [111] M. Aaboud *et al.* (ATLAS Collaboration), Search for dark matter and other new phenomena in events with an energetic jet and large missing transverse momentum using the ATLAS detector, *J. High Energy Phys.* **01** (2018) 126.
- [112] M. Aaboud *et al.* (ATLAS Collaboration), Search for long-lived particles produced in pp collisions at $\sqrt{s} = 13$ TeV that decay into displaced hadronic jets in the ATLAS muon spectrometer, *Phys. Rev. D* **99**, 052005 (2019).
- [113] G. Aad *et al.* (ATLAS Collaboration), Search for long-lived neutral particles produced in pp collisions at $\sqrt{s} = 13$ TeV decaying into displaced hadronic jets in the ATLAS inner detector and muon spectrometer, *Phys. Rev. D* **101**, 052013 (2020).
- [114] S. Schael *et al.* (ALEPH, DELPHI, L3, OPAL, SLD, LEP Electroweak Working Group, SLD Electroweak Group, SLD Heavy Flavour Group), Precision electroweak measurements on the Z resonance, *Phys. Rep.* **427**, 257 (2006).
- [115] J. Blumlein *et al.*, Limits on neutral light scalar and pseudoscalar particles in a proton beam dump experiment, *Z. Phys. C* **51**, 341 (1991).
- [116] J. Blumlein *et al.*, Limits on the mass of light (pseudo) scalar particles from Bethe-Heitler e^+e^- and $\mu^+\mu^-$ pair production in a proton—iron beam dump experiment, *Int. J. Mod. Phys. A* **07**, 3835 (1992).
- [117] J. Blumlein and J. Brunner, New exclusion limits for dark gauge forces from beam-dump data, *Phys. Lett. B* **701**, 155 (2011).
- [118] J. Blümlein and J. Brunner, New exclusion limits on dark gauge forces from proton bremsstrahlung in beam-dump data, *Phys. Lett. B* **731**, 320 (2014).
- [119] G. Bernardi, G. Carugno, J. Chauveau, F. Dicarolo, M. Dris, J. Dumarchez, M. Ferro-Luzzi, J.-M. Levy, D. Lukas, J.-M. Perreau, Y. Pons, A.-M. Touchard, and F. Vannucci, Search for neutrino decay, *Phys. Lett.* **166B**, 479 (1986).
- [120] G. Bernardi, G. Carugno, J. Chauveau, F. Dicarolo, M. Dris, J. Dumarchez, M. Ferro-Luzzi, J.-M. Levy, D. Lukas, J.-M. Perreau, Y. Pons, A.-M. Touchard, and F. Vannucci, Further limits on heavy neutrino couplings, *Phys. Lett. B* **203**, 332 (1988).
- [121] S. N. Gninenko, Stringent limits on the $\pi^0 \rightarrow \gamma X, X \rightarrow e^+e^-$ decay from neutrino experiments and constraints on new light gauge bosons, *Phys. Rev. D* **85**, 055027 (2012).
- [122] D. Gorbunov, I. Krasnov, and S. Suvorov, Constraints on light scalars from PS191 results, *Phys. Lett. B* **820**, 136524 (2021).
- [123] C. N. Yang, Selection rules for the dematerialization of a particle into two photons, *Phys. Rev.* **77**, 242 (1950).
- [124] L. D. Landau, On the angular momentum of a system of two photons, *Dokl. Akad. Nauk SSSR* **60**, 207 (1948).
- [125] J. P. Lees *et al.* (BABAR Collaboration), Search for $B \rightarrow K^{(*)}\nu\bar{\nu}$ and invisible quarkonium decays, *Phys. Rev. D* **87**, 112005 (2013).
- [126] E. Cortina Gil *et al.* (NA62 Collaboration), Measurement of the very rare $K^+ \rightarrow \pi^+\nu\bar{\nu}$ decay, *J. High Energy Phys.* **06** (2021) 093.
- [127] M. Ablikim *et al.* (BES Collaboration), Search for the Invisible Decay of J/ψ in $\psi(2S) \rightarrow \pi^+\pi^-J/\psi$, *Phys. Rev. Lett.* **100**, 192001 (2008).
- [128] F. Kling and S. Trojanowski, Forward experiment sensitivity estimator for the LHC and future hadron colliders, *Phys. Rev. D* **104**, 035012 (2021).
- [129] P. Ballett, T. Boschi, and S. Pascoli, Heavy neutral leptons from low-scale seesaws at the DUNE near detector, *J. High Energy Phys.* **03** (2020) 111.
- [130] A. Caputo, H.-T. Janka, G. Raffelt, and E. Vitagliano, Low-Energy Supernovae Severely Constrain Radiative Particle Decays, *Phys. Rev. Lett.* **128**, 221103 (2022).
- [131] J. L. Feng *et al.*, The forward physics facility at the high-luminosity LHC, *J. Phys. G* **50**, 030501 (2023).
- [132] P. Masjuan, E. R. Arriola, and W. Broniowski, Meson dominance of hadron form factors and large- N_c phenomenology, *Phys. Rev. D* **87**, 014005 (2013).
- [133] A. Ritz and S. Foroughi-Abari (private communication).
- [134] D. Gorbunov, A. Makarov, and I. Timiryasov, Decaying light particles in the SHiP experiment: Signal rate estimates for hidden photons, *Phys. Rev. D* **91**, 035027 (2015).
- [135] P. Ball and R. Zwicky, New results on $B \rightarrow \pi, K, \eta$ decay formfactors from light-cone sum rules, *Phys. Rev. D* **71**, 014015 (2005).
- [136] N. Carrasco, P. Lami, V. Lubicz, L. Riggio, S. Simula, and C. Tarantino, $K \rightarrow \pi$ semileptonic form factors with $N_f = 2 + 1 + 1$ twisted mass fermions, *Phys. Rev. D* **93**, 114512 (2016).
- [137] P. Ball and R. Zwicky, $B_{d,s} \rightarrow \rho, \omega, K^*, \phi$ decay formfactors from light-cone sum rules revisited, *Phys. Rev. D* **71**, 014029 (2005).
- [138] A. Bharucha, D. M. Straub, and R. Zwicky, $B \rightarrow V\ell^+\ell^-$ in the standard model from light-cone sum rules, *J. High Energy Phys.* **08** (2016) 098.
- [139] N. Gubernari, A. Kokulu, and D. van Dyk, $B \rightarrow P$ and $B \rightarrow V$ form factors from B -meson light-cone sum rules beyond leading twist, *J. High Energy Phys.* **01** (2019) 150.
- [140] W. G. Parrott, C. Bouchard, and C. T. H. Davies (HPQCD Collaboration), $B \rightarrow K$ and $D \rightarrow K$ form factors from fully relativistic lattice QCD, *Phys. Rev. D* **107**, 014510 (2023).
- [141] N. Gubernari, M. Reboud, D. van Dyk, and J. Virto, Dispersive analysis of $B \rightarrow K^{(*)}$ and $B_s \rightarrow \phi$ form factors, [arXiv:2305.06301](https://arxiv.org/abs/2305.06301).
- [142] T. Hahn, CUBA: A library for multidimensional numerical integration, *Comput. Phys. Commun.* **168**, 78 (2005).
- [143] D. Egana-Ugrinovic, S. Homiller, and P. Meade, Light Scalars and the Koto Anomaly, *Phys. Rev. Lett.* **124**, 191801 (2020).