

Extending the discovery potential for inelastic-dipole dark matter with FASER

Keith R. Dienes^{1,2,*}, Jonathan L. Feng^{3,†}, Max Fieg^{3,‡}, Fei Huang^{4,§}, Seung J. Lee^{5,||} and Brooks Thomas^{6,¶}

¹*Department of Physics, University of Arizona, Tucson, Arizona 85721 USA*

²*Department of Physics, University of Maryland, College Park, Maryland 20742 USA*

³*Department of Physics and Astronomy, University of California, Irvine, California 92697 USA*

⁴*Department of Particle Physics and Astrophysics, Weizmann Institute of Science, Rehovot 7610001, Israel*

⁵*Department of Physics, Korea University, Seoul 136-713, Korea*

⁶*Department of Physics, Lafayette College, Easton, Pennsylvania 18042 USA*



(Received 9 February 2023; accepted 9 May 2023; published 12 June 2023)

Neutral particles are notoriously difficult to observe through electromagnetic interactions. As a result, they naturally elude detection in most collider detectors. In this paper, we point out that neutral particles that interact through a dipole interaction can nevertheless be detected in far-forward detectors designed to search for long-lived particles (LLPs). In contrast to previous analyses that focused on neutral particles with elastic interactions, we consider inelastic interactions. This naturally leads to LLPs, and we demonstrate that FASER (and future experiments at the Forward Physics Facility) will be able to probe substantial regions of the associated parameter space. In particular, we find that FASER is capable of probing the region of parameter space wherein thermal freeze-out gives rise to an $\mathcal{O}(\text{GeV})$ dark-matter candidate with the appropriate relic abundance, as well as regions of parameter space that are difficult to probe at fixed-target experiments. FASER and its successor experiments may therefore play a critical role in the discovery of such a dark-matter candidate.

DOI: [10.1103/PhysRevD.107.115006](https://doi.org/10.1103/PhysRevD.107.115006)

I. INTRODUCTION

Since it began running almost 15 years ago, the Large Hadron Collider (LHC) has been one of the most effective tools we have for understanding the Standard Model (SM) and for probing physics beyond the Standard Model (BSM). Among its many successes, the LHC has discovered the Higgs boson and explored its properties in great detail. LHC data has also hinted at new physics in the flavor sector, and searches involving missing energy, displaced vertices, resonances, and a host of other phenomena have been used to constrain the properties of potential dark sectors [1–3]. However, despite this progress, the LHC has not yielded any definitive discovery of BSM physics.

Of course, new physics might be hiding within experimental blind spots. In particular, new physics might be found in the large-pseudorapidity regions, where there is a large flux of SM particles. The ForWard Search Experiment (FASER) [4–6] provides an ideal opportunity for probing this region. FASER is a cylindrical detector placed 476 m from the ATLAS interaction point (IP) along the beam axis. As such, FASER aims to capture the decay products of long-lived particles (LLPs). Shielded by 100 m of rock and concrete, FASER typically exploits the large SM activity near the IP as a means of dark-state production through meson decay. The FASER detector is already taking data, and an upgraded detector, known as FASER2, has been proposed, which would enhance the sensitivity of the current detector during the high-luminosity LHC (HL-LHC) era. This upgraded detector would be housed, along with several other experiments, within the proposed Forward Physics Facility (FPF) [7–9].

A common benchmark signature that FASER is well suited to probe is that of a dark photon decaying visibly to an e^+e^- pair. Such a signature can be used to constrain potential kinetic-mixing couplings of magnitudes $\epsilon \sim 10^{-5} - 10^{-4}$ with dark-photon masses $m_{A'} \lesssim 1 \text{ GeV}$ [4]. In addition to dark-photon models, FASER and the FPF have discovery prospects for a variety of new-physics scenarios in which an LLP, such as an axion-like particle, dark scalar,

*dienes@arizona.edu

†jlf@uci.edu

‡mfieg@uci.edu

§fei.huang@weizmann.ac.il

||sjlee@korea.ac.kr

¶thomasbd@lafayette.edu

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

or heavy neutral lepton, decays to a pair of detectable SM particles. Although such scenarios are well studied in the literature (see, e.g., Ref. [5] and references therein), it is important to fully explore additional signatures to which FASER might also be sensitive.

In this paper, we examine the detection prospects for one such LLP signature at FASER—a monophoton signature that arises in scenarios in which the photon field couples to a pair of neutral BSM particles χ_0 and χ_1 with different masses $m_1 > m_0$ via an effective magnetic-dipole-moment (MDM) or electric-dipole-moment (EDM) operator. Such a scenario is particularly interesting from a phenomenological perspective, in part because it arises in the context of inelastic-dark-matter scenarios [10–12]. The inelastic MDM operator was also recently studied in the context of sterile neutrinos [13]. If χ_0 and χ_1 are sufficiently light, they can be produced together at a significant rate at the LHC via meson-decay processes involving such MDM and EDM operators. Moreover, these operators render χ_1 unstable, since χ_1 can decay via the process $\chi_1 \rightarrow \chi_0 \gamma$. Thus, if the resulting χ_1 decay width is of the appropriate order, particles of this species emitted in the forward direction at the LHC could potentially decay inside the FASER detector. The calorimeter of the FASER detector can then detect the resulting monophoton signal [14,15].

From a theoretical perspective, inelastic-dark-matter models of this sort are also of interest because they represent the simplest realization of a nonminimal dark sector involving N dark states χ_i with similar quantum numbers, where $i = 0, 1, \dots, N - 1$. The phenomenology of such dark sectors can differ significantly from that of minimal dark sectors, due to the presence of the additional dark-sector states (for a review, see Ref. [16]). Indeed, the presence of the additional dark-sector states generically gives rise to decay processes in which a heavier χ_i particle decays into a final state involving one or more lighter χ_i particles. Such decay processes alter the complementarity relations between different experimental probes of the dark sector [17,18]. Moreover, they can also lead to a variety of striking signatures at colliders [19,20] and even open up new ways of addressing the dark-matter problem, such as those that arise within the Dynamical Dark Matter (DDM) framework [21–23]. Nonminimal dark sectors involving multiple states with similar quantum numbers—and indeed often involving large values of N —emerge naturally in a variety of BSM contexts, including theories with extra spacetime dimensions, theories involving confining hidden-sector gauge groups, and string theory. A two-component inelastic-dark-matter model of the sort on which we shall focus in this paper may be viewed as the simplest example of a nonminimal dark sector in which these phenomenological possibilities arise.

One of the most important features of inelastic dark matter is that it provides a mechanism by which event rates at direct-detection experiments can be suppressed.

Moreover, the dynamics involved in thermal freeze-out is significantly more complicated in such models, with annihilation, co-annihilation, up- and down-scattering, and decay processes all playing a role in generating the relic abundance of the lighter dark-sector state. Realizations of inelastic dark matter in which the dark-sector states couple to the visible sector via dipole-moment interactions in particular can give rise to a dark-matter abundance in accord with observation without running afoul of the applicable constraints [24–26]. Indeed, data from a variety of sources can constrain the parameter space of inelastic-dipole-dark-matter models, including data from collider experiments, fixed-target experiments, and gamma-ray telescopes. That said, there are also interesting regions of this parameter space in which existing searches have little or no sensitivity. Notably these include those in which the spectra are highly compressed and the lighter dark-sector state accounts for the entirety of the observed dark-matter abundance.

In this paper, we investigate the extent to which searches for a monophoton signature of $\chi_1 \rightarrow \chi_0 \gamma$ decays at FASER and FASER2 can constrain the parameter space of such inelastic-dipole-dark-matter models. As we shall see, these detectors provide coverage of this parameter space within certain regions that are not well constrained by existing searches. Indeed, FASER and FASER2 can probe parameter-space regions within which the splitting between the masses m_1 and m_0 of the two dark-sector states χ_1 and χ_0 is roughly $\Delta m \equiv m_1 - m_0 \sim \mathcal{O}(10^{-2})m_0$ and where $m_0 \lesssim 5$ GeV. Remarkably, the reach of FASER and FASER2 within that parameter space includes regions wherein the relic abundance of χ_0 obtained from thermal freeze-out agrees with observed abundance of dark matter. We note also that the small mass-splittings Δm that make the heavier state χ_1 long-lived also suppress the energy of the photon into which it decays, thereby rendering the resulting monophoton signal challenging to detect at fixed-target experiments. At FASER, however, this suppression is compensated by the significant boosts provided to these decaying particles as a result of the high collision energies achieved at the LHC. Thus, as we shall see, FASER is capable of probing regions of parameter space that are difficult to probe at fixed-target experiments.

This paper is organized as follows. In Sec. II, we review the manner in which the dark-sector states couple to the visible sector in inelastic-dipole-dark-matter models. In Sec. III, we calculate the event rates for the signal and background processes relevant for monophoton searches at FASER and FASER2 within the context of these models and assess the discovery reach of these detectors within the parameter space of these models. In Sec. IV, we examine the annihilation and co-annihilation processes involved in the thermal freeze-out of χ_0 and χ_1 in the early universe and identify the regions of parameter space within which relic χ_0 particles constitute the majority of the dark matter at the

present time. In Sec. V, we examine the additional considerations that constrain the parameter space of inelastic-dipole-dark-matter models in the region where FASER is sensitive. In Sec. VI, we conclude by summarizing our results and discussing possible directions for future work. Details of the calculation of meson branching fractions and thermal relic densities, along with a brief discussion of the detection prospects at SHiP [27,28], a proposed experiment complementary to FASER which will probe comparable lifetimes, are given in the appendices.

II. INELASTIC-DIPOLE DARK MATTER

We consider an inelastic-dark-matter model in which the two dark-sector particles χ_0 and χ_1 are fermions that couple to the visible sector via an effective MDM or EDM operator in the interaction Lagrangian of the form

$$\begin{aligned}\mathcal{O}_m &= \frac{1}{\Lambda_m} \bar{\chi}_1 \sigma^{\mu\nu} \chi_0 F_{\mu\nu} + \text{H.c.} \\ \mathcal{O}_e &= \frac{1}{\Lambda_e} \bar{\chi}_1 \sigma^{\mu\nu} \gamma^5 \chi_0 F_{\mu\nu} + \text{H.c.},\end{aligned}\quad (2.1)$$

where $F_{\mu\nu}$ is the field-strength tensor for the photon field and where $\Lambda_{\mathcal{O}}$, with $\mathcal{O} = \{m, e\}$, denotes the corresponding operator-suppression scale.¹ We note that such effective operators can arise in a variety of UV-physics scenarios. For example, they can arise as a consequence of loop-level processes involving heavy particles with masses of $\mathcal{O}(\text{TeV})$ or higher that carry SM charges. At scales above the electroweak-symmetry-breaking scale, these operators would presumably be replaced by operators with a similar structure, but with the $U(1)_Y$ and/or $SU(2)_L$ gauge bosons present in the unbroken phase of the SM in place of the photon [25]. However, this phase of the SM is not relevant for the phenomenological implications of inelastic-dipole dark matter that we consider in this paper.

We note that these operators can also be viewed as arising from terms in the elastic tensor and pseudotensor currents $J^{\mu\nu} = \bar{\chi} \sigma^{\mu\nu} \chi$ and $J^{\mu 5} = \bar{\chi} \sigma^{\mu\nu} \gamma^5 \chi$ involving a Dirac fermion χ that also acquires a Majorana mass [24–26]. Such a Majorana mass gives rise to a mass-splitting between the fermion mass-eigenstates—states that play the role of our χ_0 and χ_1 —and to operators with the inelastic coupling structures given in Eq. (2.1). We parametrize this mass-splitting in terms of the dimensionless ratio

$$\Delta \equiv \frac{\Delta m}{m_0} \equiv \frac{m_1 - m_0}{m_0}. \quad (2.2)$$

¹An alternative parametrization is often used in the literature in which the strengths of the effective MDM and EDM interactions are parametrized by the dimensionful coupling coefficients μ_χ and d_χ , which are related to our operator-suppression scales Λ_m and Λ_e by $\mu_\chi = 2/\Lambda_m$ and $d_\chi = 2/\Lambda_e$.

Since Dirac and Majorana terms in the dark-fermion mass matrix need not necessarily have a common origin, Δ can in principle take a broad range of values. As such, we take the parameter space of our inelastic-dipole-dark-matter model to be characterized by the three parameters m_0 , Δ , and $\Lambda_{\mathcal{O}}$.

The interactions in Eq. (2.1) have a variety of phenomenological implications. For example, in the presence of such interactions, χ_0 and χ_1 particles can be produced both at colliders and at beam-dump experiments via processes of the form $f\bar{f} \rightarrow \chi_0\bar{\chi}_1$, where f denotes a SM fermion, via bremsstrahlung-like processes involving a virtual photon, or via the decays of vector and pseudoscalar mesons. Once produced, these particles can be detected in a variety of ways due to the up-scattering, down-scattering, and decay processes to which these same interactions give rise. Indeed, potential experimental signatures of inelastic-dipole dark matter include the scattering of χ_0 or χ_1 with atomic nuclei, events involving sizable missing transverse energy at colliders, and rare meson decays [29–31]. Existing limits on these processes place nontrivial constraints on our parameter space.

The decay process $\chi_1 \rightarrow \chi_0\gamma$ provides an additional experimental probe of inelastic-dipole-dark-matter models. Indeed, searches for decays of this sort at B -factories and at the main LHC detectors likewise place nontrivial constraints on our parameter space [25]. However, as discussed in the Introduction, this same decay process can also potentially lead to signatures at FASER and FASER2 within regions of that parameter space that are largely unconstrained. One particularly interesting such region of parameter space turns out to be that in which $\Delta \sim \mathcal{O}(0.01)$ and $\mathcal{O}(\text{MeV}) < m_0 < \mathcal{O}(\text{GeV})$. Within this region, not only can a significant flux of χ_i pairs be generated in the forward direction, but the lifetime τ of χ_1 is also such that a significant fraction of χ_1 decays occur within the FASER detector. Indeed, as we shall justify below, the lab-frame decay length $\bar{d} = \gamma\beta\tau$, where β and γ are the usual relativistic factors, can be expressed as a product of factors of the form

$$\begin{aligned}\bar{d} &\approx (600 \text{ m}) \times \left(\frac{500 \text{ MeV}}{m_0}\right)^4 \left(\frac{0.01}{\Delta}\right)^3 \\ &\times \left(\frac{\Lambda_{\mathcal{O}}}{15.5 \text{ TeV}}\right)^2 \left(\frac{E_{\chi_1}}{\text{TeV}}\right),\end{aligned}\quad (2.3)$$

where E_{χ_1} is the energy of χ_1 . Given that FASER is located approximately 476 m away from the ATLAS detector, this expression for \bar{d} makes it clear that the region of parameter space within which $\Delta \sim \mathcal{O}(0.01)$ and $\mathcal{O}(\text{MeV}) < m_0 < \mathcal{O}(\text{GeV})$, while $\Lambda_{\mathcal{O}}$ is near the TeV scale, is one in which the prospects of observing χ_1 decays insider FASER are particularly auspicious. Since the proposed location of FASER2 within the FPF lies a similar distance away from the ATLAS detector, this region of

parameter space is auspicious for detection at FASER2 as well.

As we shall see, this region of inelastic-dipole-dark-matter parameter space is interesting not only because it is auspicious for detection at FASER; it is also interesting from a dark-matter perspective. Indeed, within this region the abundance of χ_0 obtained from thermal freeze-out and the subsequent decays of χ_1 can accord with the dark-matter abundance inferred from Planck data [32]. Moreover, since the value of the mass-splitting $\Delta m \gtrsim 10$ keV within this parameter-space region of interest lies well above the lab-frame kinetic energy of a typical χ_0 particle of mass $\mathcal{O}(\text{MeV}) < m_0 < \mathcal{O}(\text{GeV})$ in the Milky-Way halo, event rates at direct-detection experiments are highly suppressed. As a result, such experiments are insensitive to dipole-interacting dark matter within this region of parameter space.

III. SIGNAL AT FASER

A. Production

A number of processes contribute to the overall production rate for χ_1 particles in the far-forward direction at the LHC in inelastic-dipole-dark-matter models. These include production from meson decay, Drell-Yan production, and secondary production—i.e., the production of χ_1 particles due to the up-scattering of χ_0 particles produced at the IP as they travel through the intervening material between the IP and the FASER detector [33].

The decays of neutral mesons produced at the ATLAS IP turn out to provide the dominant contribution to this overall production rate. Indeed both vector and pseudoscalar mesons are produced at the LHC in copious amounts in the forward direction towards the FASER detector, which is coaxial with the ATLAS experiment. Moreover, the decays of both kinds of meson contribute to χ_1 production. Within the regime in which $\mathcal{O}(\text{MeV}) \lesssim (2 + \Delta)m_0 \lesssim \mathcal{O}(\text{GeV})$, the relevant pseudoscalar mesons are π^0 , η , and η' , while the relevant vector mesons are ρ , ω , ϕ , J/ψ , $\psi(2S)$, and $\Upsilon(nS)$. The dominant contribution to χ_1 production from each of these meson species M is due to the decays of mesons that are produced in the far-forward direction—i.e., with pseudorapidities $\eta > \eta_F$, where $\eta_F = 9.2$ and $\eta_F = 7.2$ for FASER and FASER2, respectively—and have lab-frame energies $E_M \sim \mathcal{O}(\text{TeV})$. Mesons with such energies are sufficiently highly boosted that the vast majority of χ_1 particles produced by their decays will likewise travel in the far-forward direction, regardless of the angle that the momentum vector of the χ_1 makes with the momentum vector of the decaying meson in the rest frame of the meson. Thus, we may obtain a rough estimate of the total cross-section $\sigma_{\chi_1 M}^{(\text{eff})}$ for the production of χ_1 particles in the direction of FASER from the decay of a particular meson species M simply by multiplying the production cross-section σ_M for mesons

of that species with $\eta \geq \eta_F$ by the overall branching fraction for M into final states including χ_1 . We obtain the σ_M for each relevant meson species from the FORESEE code package [34].

A pseudoscalar meson P produces χ_i particles primarily through the three-body decay process $P \rightarrow \gamma\gamma^* \rightarrow \bar{\chi}_0\chi_1\gamma$. One could also consider the two-body decay $P \rightarrow \bar{\chi}_0\chi_1$, but the branching fraction for this process is loop-suppressed (because the leading diagram arises at one-loop order), helicity-suppressed, and further suppressed by $\Lambda_{\mathcal{O}}^{-4} \sim (10 \text{ TeV})^{-4}$, so we shall not consider this production process further. By contrast, a vector meson V produces these particles primarily through the two-body decay process $V \rightarrow \bar{\chi}_0\chi_1$. The corresponding decay processes for the case of an elastic MDM or EDM interaction were considered in Refs. [29–31]. Full expressions for the branching fractions $\text{BR}(V \rightarrow \bar{\chi}_0\chi_1)$ and $\text{BR}(P \rightarrow \bar{\chi}_0\chi_1\gamma)$ for the aforementioned vector- and pseudoscalar-meson decay processes are provided in Appendix A for both the MDM- and an EDM-interaction cases. We note that both $\text{BR}(V \rightarrow \bar{\chi}_0\chi_1)$ and $\text{BR}(P \rightarrow \bar{\chi}_0\chi_1\gamma)$ are, to a good approximation, proportional to the square of the meson mass m_M in the $(2 + \Delta)m_0 \ll m_M$ regime. Thus, although the cross-sections for the production of light pseudoscalars such as the π^0 in the forward direction at the LHC are significantly higher than those for much heavier vector mesons, the likelihood that each such vector meson will produce a χ_1 when it decays is enhanced simply by virtue of its being heavier. Moreover, the $\text{BR}(P \rightarrow \bar{\chi}_0\chi_1\gamma)$ for all pseudoscalar mesons, regardless of their mass, experience an additional suppression relative to the $\text{BR}(V \rightarrow \bar{\chi}_0\chi_1)$ for vector mesons due to phase-space considerations. As a consequence of these two effects, we find that the decays of vector mesons in fact provide the dominant contribution to the production rate of χ_1 particles in the forward direction at the LHC within our parameter-space region of interest.

Given that vector-meson decays to χ_i pairs dominate the production rate of χ_1 particles in the forward direction, it is instructive to compare the branching-fraction expressions $\text{BR}_{\text{MDM}}(V \rightarrow \bar{\chi}_0\chi_1)$ and $\text{BR}_{\text{EDM}}(V \rightarrow \bar{\chi}_0\chi_1)$ for these decays in the MDM- and EDM-interaction cases, respectively. In the regime in which $m_V \gg (2 + \Delta)m_0$, these two expressions coincide and are well approximated by the single expression

$$\text{BR}(V \rightarrow \bar{\chi}_0\chi_1) \approx \frac{m_V^2}{8\pi\alpha\Lambda_{\mathcal{O}}^2} \text{BR}(V \rightarrow e^+e^-), \quad (3.1)$$

where α is the fine-structure constant. By contrast, as the dark-fermion masses are increased to the point at which $(2 + \Delta)m_0$ becomes comparable with m_V , the expressions for $\text{BR}_{\text{MDM}}(V \rightarrow \bar{\chi}_0\chi_1)$ and $\text{BR}_{\text{EDM}}(V \rightarrow \bar{\chi}_0\chi_1)$ begin to deviate from each other appreciably. For example, in the elastic limit—i.e., the limit in which $\Delta \rightarrow 0$ —the ratio of

these two branching-fraction expressions takes the particularly simple form

$$\frac{\text{BR}_{\text{MDM}}(V \rightarrow \bar{\chi}_0 \chi_1)}{\text{BR}_{\text{EDM}}(V \rightarrow \bar{\chi}_0 \chi_1)} \xrightarrow{\Delta \rightarrow 0} \frac{\Lambda_e^2}{\Lambda_m^2} \left(\frac{M_V^2 + 8m_0^2}{M_V^2 - 4m_0^2} \right). \quad (3.2)$$

Thus, in this limit, the production rate of χ_i pairs from vector-meson decay is larger for an MDM interaction than it is for an EDM interaction with the same value of $\Lambda_{\mathcal{O}}$. We find that this qualitative result likewise holds in the inelastic case in which $\Delta \neq 0$ as well. This can be understood as a consequence of angular-momentum conservation. In the decaying meson's center-of-mass frame, the MDM operator creates the $\bar{\chi}_0 \chi_1$ final state in the $|SS_z\rangle = |10\rangle$ spin state, and since angular-momentum conservation requires $J = 1$, the final-state orbital angular momentum is $L = 0$. By contrast, the EDM operator creates a final state with spin state $|SS_z\rangle = |00\rangle$ and $L = 1$. As a result, the partial width in the EDM case is P -wave suppressed, as can be seen in the squared amplitude in Eq. (A4), and thus FASER will have greater sensitivity for large m_0 for the MDM case.

In Fig. 1, we show the effective cross-sections $\sigma_{\chi_1 M}^{(\text{eff})}$ for the production of χ_1 particles in the direction of FASER from the decay of several relevant meson species M as functions of m_0 for the choice of $\Delta = 0.01$. In order to display our results in as model-independent a manner as possible, we scale each effective cross-section by a factor of $(\Lambda_{\mathcal{O}}/\text{GeV})^2$ to cancel the overall factor of $\Lambda_{\mathcal{O}}^{-2}$ common to all of the $\sigma_{\chi_1 M}^{(\text{eff})}$. We observe that for this value of Δ , the decays of vector mesons collectively provide the dominant

contribution to the production rate of χ_1 particles in the direction of FASER over the entire range of m_0 shown. Indeed, even for values of m_0 sufficiently small that the decay process $\pi^0 \rightarrow \gamma \bar{\chi}_0 \chi_1$ is kinematically accessible, we find that the contribution to the total production rate from π^0 decay is subleading in comparison with the contributions from the decays of the heavier vector mesons ω and ϕ . Moreover, we note that these results are not particularly sensitive to the value of Δ , provided that $\Delta \ll 1$.

In addition to the contribution from meson decays, a contribution to the overall production rate of χ_1 particles at the LHC also arises from Drell-Yan processes of the form $q\bar{q} \rightarrow \bar{\chi}_1 \chi_0$. Since kinematic considerations imply that only mesons M with masses $m_M > (2 + \Delta)m_0$ contribute to the production of χ_1 particles through their decays, this Drell-Yan contribution can in principle be important in the regime in which $(2 + \Delta)m_0$ is large. In order to determine the size of this latter contribution, we evaluate the differential cross-section $d^2\sigma/(dE_{\chi_1} d\eta_{\chi_1})$ for the Drell-Yan production of χ_1 particles as a function of their energy E_{χ_1} and pseudorapidity η_{χ_1} using the MG5_AMC@NLO [35] code package. We focus on the regime in which $(2 + \Delta)m_0 > 1$ GeV, as parton-distribution functions are not well defined for momentum transfers below this scale. Based on the results of these simulations, we find that the contribution to the production rate of χ_1 particles from Drell-Yan processes is subleading in comparison with the contribution from meson decay.

Yet another contribution to the production rate for χ_1 arises due to secondary production [33]. Using MG5_AMC@NLO we evaluate this contribution by




FIG. 1. Effective cross-sections $\sigma_{\chi_1 M}^{(\text{eff})}$ for the production of χ_1 particles in the direction of FASER from the decay of several relevant meson species M , shown as functions of m_0 for $\eta_F = 9.2$. In order to display our results in as model-independent a manner as possible, we scale each cross-section by a factor of $(\Lambda_{\mathcal{O}}/\text{GeV})^2$ to cancel the overall factor of $\Lambda_{\mathcal{O}}^{-2}$ common to all of the $\sigma_{\chi_1 M}^{(\text{eff})}$. The solid curves correspond to the case of a MDM operator, while the dashed curves correspond to the case of an EDM operator. The results shown here correspond to the case in which $\Delta = 0.01$; however, we note that these results are not particularly sensitive to the value of Δ , provided that $\Delta \ll 1$.

calculating the phase-space distribution of χ_1 after up-scattering via $\chi_0 + A \rightarrow \chi_1 + A$ off a nucleus A in the material located between the ATLAS IP and FASER. We incorporate appropriate form factors for the concrete, rock, and for the FASER ν tungsten target and calculate the event rate for subsequent χ_1 decay in FASER [33]. The region of parameter space for which this effect becomes potentially important is that within which $10^1 \text{ GeV} \lesssim \Lambda_{\mathcal{O}} \lesssim 10^3 \text{ GeV}$ and $m_0 \sim \mathcal{O}(100) \text{ MeV}$. In particular we find that within this region of parameter space secondary production off tungsten in FASER ν yields a potentially significant contribution to the signal rate. By contrast, we find that secondary production off the concrete and rock located between the ATLAS IP and FASER yields a subleading contribution to the signal rate.

While secondary production can indeed be important within this region of parameter space, this region turns out to be excluded by data from *BABAR*, *CHARM-II*, and *LEP*, as we shall discuss in detail in Sec. V. By contrast within regions of parameter space not excluded by current data, we find that the contribution to the total event rate is typically well below 1% of the contribution from meson decay.

B. Signal rate

After a χ_1 particle is produced at ATLAS in our inelastic-dipole-dark-matter model, it decays via the process $\chi_1 \rightarrow \chi_0 \gamma$. The lab-frame energy of the photon produced by the decay of a χ_1 particle with lab-frame energy E_{χ_1} is

$$E_\gamma = \frac{\Delta(2 + \Delta)}{2(1 + \Delta)} [\gamma + (\gamma^2 - 1)^{1/2} \cos \theta_{\gamma\chi_1}] m_0, \quad (3.3)$$

where $\gamma \equiv E_{\chi_1}/m_1$ is the usual relativistic boost factor and where $\theta_{\gamma\chi_1}$ is the angle in the rest frame of the χ_1 particle between the momentum vector of the photon and the axis along which the χ_1 particle is traveling. Since the χ_1 particle effectively inherits its boost factor from the decaying meson that produced it, and since the characteristic energy scale for the mesons is $E_M \sim \mathcal{O}(1 \text{ TeV})$, the characteristic size for this boost factor is $\gamma \sim E_M/m_M \sim \mathcal{O}(10^3)$. Thus, the characteristic photon energy for these decays within our parameter space is $E_\gamma \approx \gamma m_0 \Delta \sim \mathcal{O}(10 \text{ GeV})$. A single photon with such an energy can be detected by the calorimeter in either FASER or FASER2 [6].

The probability of observing a statistically significant number of monophoton events at FASER during Run 3 or at FASER2 during the HL-LHC era depends both on the geometry of the corresponding detector and on the decay properties of χ_1 itself. FASER and FASER2 each have a cylindrical geometry in which the cylinder is coaxial with the beam-collision axis at the ATLAS IP. The components of FASER and FASER2 each include, starting from the side closest to the ATLAS detector, a dedicated decay volume, a tracking volume interleaved with a series of

TABLE I. The parameters that characterize the detector geometry of FASER and FASER2, along with the integrated luminosity expected at ATLAS during the corresponding running period for each detector.

| Detector | h | R | L | η_F | \mathcal{L}_{int} |
|----------|-------|-------|-------|----------|----------------------------|
| FASER | 3.5 m | 10 cm | 476 m | 9.2 | 300 fb $^{-1}$ |
| FASER2 | 20 m | 1 m | 620 m | 7.2 | 3 ab $^{-1}$ |

charged-particle trackers, a ‘‘preshower’’ station, and a calorimeter. Further details about each component are provided in Ref. [6]. The radius R of the cylindrical decay and tracking volumes, their combined length h along the axial direction of the cylinder, the distance L between the side of the decay volume closest to the ATLAS detector and the ATLAS IP, the value of η_F that follows from the values of L and h , and the integrated luminosity \mathcal{L}_{int} expected at ATLAS during the corresponding running period are provided in Table I.

We note that the values of h given in Table I include both the dedicated decay volume of the detector and the tracking volume. Indeed, the detection of a single photon within the calorimeter of FASER or FASER2 does not rely on information from the tracker. Moreover, the scintillators separating the decay volume and tracking chambers of either detector are largely transparent to photons with $E_\gamma \sim \mathcal{O}(10 \text{ GeV})$ energies. Taken together, these two considerations imply that the tracking volume can also be viewed as an extension of the decay volume in searches for a monophoton signal of $\chi_1 \rightarrow \chi_0 \gamma$ decays within these detectors.

The characteristic decay length of χ_1 in the lab frame is

$$\bar{d} = \frac{|\vec{\mathbf{p}}_{\chi_1}|}{m_1 \Gamma_{\chi_1}}, \quad (3.4)$$

where $\vec{\mathbf{p}}_{\chi_1}$ and Γ_{χ_1} denote the lab-frame momentum and proper decay width of χ_1 , respectively. Since all of the meson species that contribute significantly to the production rate of χ_1 particles decay promptly within the ATLAS detector, \bar{d} represents the characteristic distance away from the IP that χ_1 particles travel before they decay. Within our parameter-space regime the proper decay width of χ_1 is well approximated by

$$\Gamma_{\chi_1} = \frac{1}{2\pi\Lambda_{\mathcal{O}}^2} \left(\frac{m_1^2 - m_0^2}{m_1} \right)^3 \approx \frac{4m_0^3 \Delta^3}{\pi\Lambda_{\mathcal{O}}^2} \quad (3.5)$$

in both the MDM and EDM cases.

Since $R \ll h$ for both FASER and FASER2, we may approximate the distance that a χ_1 particle would traverse inside the decay volume of the detector if it did not decay as h , regardless of its pseudorapidity η_{χ_1} . Thus, the probability

that a given χ_1 will decay inside FASER or FASER2 is approximately

$$P(\bar{d}, h, L) \approx e^{-L/\bar{d}}(1 - e^{-h/\bar{d}}). \quad (3.6)$$

The total expected number of signal events from χ_1 decays within the FASER or FASER2 detector can therefore be approximated as

$$N_{\chi_1} \approx \mathcal{L}_{\text{int}} \int_{\eta_F}^{\infty} d\eta_{\chi_1} \int_{m_1}^{\infty} dE_{\chi_1} \left[\sum_M \frac{d\sigma_{\chi_1 M}^{\text{(eff)}}}{dE_{\chi_1} d\eta_{\chi_1}} \times P(\bar{d}(E_{\chi_1}), h, L) \right], \quad (3.7)$$

where $d\sigma_{\chi_1 M}^{\text{(eff)}}/(dE_{\chi_1} d\eta_{\chi_1})$ is the differential cross-section for the production of χ_1 particles as a function of E_{χ_1} and η_{χ_1} . We evaluate each of these differential cross-sections numerically by integrating the product of the differential production cross-section for M obtained from the FORESEE package [34] and the appropriate phase-space factor over possible meson momenta.

C. Background

The primary detector backgrounds for an $E_\gamma \sim \mathcal{O}(\text{GeV})$ monophoton signal of inelastic-dipole dark matter at FASER or FASER2 are those associated with muons or neutrinos produced at the ATLAS IP—or with muons produced by cosmic-ray showers—interacting with the calorimeter material [14,15,36].

Two scintillator stations located in front of the decay volume of FASER provide a veto on muons and other charged particles emanating from the direction of the ATLAS IP, the efficiency of which is around 99.99% for each station individually [6]. As a result, the contribution to the overall background-event rate from such muons is expected to be negligible. An additional contribution to this overall rate also arises from events in which a muon is produced at the IP with a pseudorapidity η sufficiently large that it travels in the general direction of FASER, but sufficiently small that its path does not intersect either of the scintillator stations. If such a muon interacts with the external detector apparatus itself or some part of the surrounding infrastructure, this interaction can nevertheless produce a shower of particles. Depending on the directions in which these particles are emitted, some of these particles could potentially reach and deposit their energies in the calorimeter [37]. However, this contribution is also expected to be negligible, given that the particles produced in the shower are typically emitted in a narrow cone around the axis along which the muon was originally traveling.

The majority of cosmic-ray muons that pass through the FASER detector would likewise evade the scintillator veto. However, these muons also behave as minimum ionizing

particles and typically give rise to energy deposits in the FASER calorimeter of $\mathcal{O}(\text{MeV})$ —an energy scale far below the energy scale associated with the signal process. Moreover, the contribution to the background-event rate from cosmic-ray muons can be further reduced by correlating timing information for calorimeter events at FASER with timing information for collision events at ATLAS. Given these considerations, this contribution is also expected to be negligible.

Finally, neutrinos produced in the far-forward direction at the ATLAS IP that undergo deep inelastic scattering with the material in the FASER calorimeter can give rise to energy deposits similar to those associated with the signal process. Background events of this sort can be distinguished from signal events on the basis of information from the high-granularity preshower detector [36,38]. Nevertheless, a residual contribution to the background-event rate still arises from two classes of neutrino events. The first class consists of events in which the scattering of the neutrino with the calorimeter generates a backsplash involving one or more photons that are subsequently detected by the preshower detector. We can significantly reduce the contribution to the background-event rate from the first class of events by imposing a maximum cut on the total energy deposited in the calorimeter, which is far higher—typically around $\mathcal{O}(100 \text{ GeV})$ [39]—for these neutrino-background events than it is for the events that constitute our signal. While we do not impose such a cut explicitly in our analysis, we note that the rate of these backsplash photon events is not expected to be high, even at energies above $\mathcal{O}(100 \text{ GeV})$. Thus, we do not expect the imposition of such a cut to have a significant impact on our results.

The second class of background events consists of those in which a neutrino scatters in or near the final tracking layer and produces charged particles that are not detected by any of the trackers, but are nevertheless detected by the calorimeter. While a precise determination of this background would require a full detector simulation, the event rate for neutrino events of this class is expected to be extremely small. Nevertheless, a TeV-neutrino scattering event is expected to produce $\mathcal{O}(10)$ charged particles with larger energies than our photon signal. Event-selection criteria designed to reject events which contain a significant number of such particles can therefore be effective in reducing this background.

Photons from $\chi_1 \rightarrow \chi_0 \gamma$ decay typically have energies $E_\gamma \sim (1 \text{ TeV})\Delta$ within our parameter-space region of interest. However, the FASER and FASER2 detectors are primarily designed to look for $\mathcal{O}(\text{TeV})$ energy deposits, and the detection efficiency $\epsilon_\gamma(E_\gamma)$ for a monophoton signal with $E_\gamma \ll \mathcal{O}(\text{TeV})$ as a function of E_γ is not well established. In order to illustrate the impact of this efficiency on the sensitivity of these detectors to our signal process, we focus for simplicity on the case in which $\epsilon_\gamma(E_\gamma)$




FIG. 2. Contours within the $(\Delta, E_{\gamma,\min})$ -plane of the fraction $f_{\gamma}(E_{\gamma,\min})$ of signal events that would be detected at FASER for a given photon-energy threshold $E_{\gamma,\min}$, relative to the number that would be detected for $E_{\gamma,\min} = 0$. The solid, dashed, and dotted contours appearing in each panel correspond to different combinations of the parameters m_0 and $\Lambda_{\mathcal{O}}^{-1}$. These combinations are specified in the legend using the notation $(m_0, \Lambda_{\mathcal{O}}^{-1})$, where m_0 and $\Lambda_{\mathcal{O}}^{-1}$ are given in units of MeV and GeV^{-1} , respectively. Two such contours are included for each parameter combination. The upper contour corresponds to $f_{\gamma}(E_{\gamma,\min}) = 0.1$, while the lower contour corresponds to $f_{\gamma}(E_{\gamma,\min}) = 0.5$. The three panels of the figure simply display different regions of the $(\Delta, E_{\gamma,\min})$ -plane in order to illustrate the shapes of these contours for the three m_0 and $\Lambda_{\mathcal{O}}$ combinations shown therein.

can be modeled in terms of a binary detection threshold $E_{\gamma,\min}$ —i.e., we take $\epsilon_{\gamma}(E_{\gamma}) = \Theta(E_{\gamma} - E_{\gamma,\min})$. In Fig. 2, we display contours within the $(\Delta, E_{\gamma,\min})$ -plane of the fraction $f_{\gamma}(E_{\gamma,\min})$ of signal events that would be detected at FASER for given $E_{\gamma,\min}$, relative to the number that would be obtained for $E_{\gamma,\min} = 0$. The solid, dashed, and dotted contours appearing in each panel of the figure correspond to different combinations of the parameters m_0 and $\Lambda_{\mathcal{O}}$. Two such contours—one corresponding to $f_{\gamma}(E_{\gamma,\min}) = 0.1$ and the other to $f_{\gamma}(E_{\gamma,\min}) = 0.5$ —are included for each parameter combination. The three panels of the figure simply display different regions of the $(\Delta, E_{\gamma,\min})$ -plane in order to illustrate the shape of these contours for the three m_0 and $\Lambda_{\mathcal{O}}$ combinations shown therein. The results displayed in the figure correspond to the case of a MDM interaction. However, the value of $f_{\gamma}(E_{\gamma,\min})$ for an MDM interaction does not differ by more than 5% of the corresponding value obtained for an EDM interaction along the contours shown. In obtaining the results shown in Fig. 3, we have employed a cut on the photon energy of $E_{\gamma} > 300$ MeV, which is well above expected energy deposits coming from minimum ionizing particles.

D. Results

In Fig. 3, we show contours of the integrated luminosity required to observe a monophoton signal of inelastic-dipole dark matter at FASER (red curves) and FASER2 (blue curves) within the $(\Lambda_{\mathcal{O}}^{-1}, m_0)$ -plane for a few of different choices of Δ . Since the expected number of background events is effectively zero, as discussed in Sec. III C, we adopt the observation of at least three signal events as our discovery criterion. We emphasize that the luminosity contours can also be interpreted as event contours for a fixed total luminosity, \mathcal{L}_{int} (cf. Table I), i.e. for a contour,

the number of events can be read as $N_{\chi_1} = 3 \times (\mathcal{L}_{\text{int}}/\mathcal{L})$, where \mathcal{L} is the luminosity of the contour. This can be useful in the case where there are unexpected backgrounds. The top, middle, and bottom panels in each column correspond to the choices $\Delta = 0.001$, $\Delta = 0.01$, and $\Delta = 0.05$, respectively. For smaller Δ , we expect to be sensitive only to values of $\Lambda_{\mathcal{O}}^{-1}$ which lie within the region excluded by constraints from CHARM-II and LEP—constraints which are insensitive to Δ for $\Delta \ll 1$. While we do not consider smaller Δ for this reason, we note that we expect FASER also to have some sensitivity for smaller Δ . The results in the left column correspond to the case of a MDM interaction, while the results in the right column correspond to the case of an EDM interaction. The dashed black curve in each panel indicates the contour along which the present-day relic abundance of χ_0 —a derivation of which shall be presented in Sec. IV—accords with the dark-matter abundance inferred from Planck data [32]. For $\Delta \ll 1$, one expects our model to approach the elastic limit (i.e. $\Delta = 0$), and thus the relic target should remain fixed for $\Delta < 10^{-3}$. The gray regions in each panel are excluded by existing constraints—constraints that will be discussed in more detail in Sec. V.

The overall shape of the sensitivity contours that appear in Fig. 3 can be understood as follows. For a given integrated luminosity and a given combination of m_0 and Δ , the range of $\Lambda_{\mathcal{O}}$ for which a statistically significant number of χ_1 particles decay within the FASER or FASER2 detector is bounded from both above and below. The lower bound stems from the requirement that \bar{d} be sufficiently large that a non-negligible number of χ_1 particles reach the decay volume before they decay. Taken together, Eqs. (3.4) and (3.5) imply that $\bar{d} \propto \Lambda_{\mathcal{O}}^2/(m_0^4 \Delta^3)$ for fixed $|\vec{p}_{\chi_1}|$. This implies that for a fixed value of Δ , we expect the contour in the $(\Lambda_{\mathcal{O}}^{-1}, m_0)$ -plane above which a significant number of χ_1

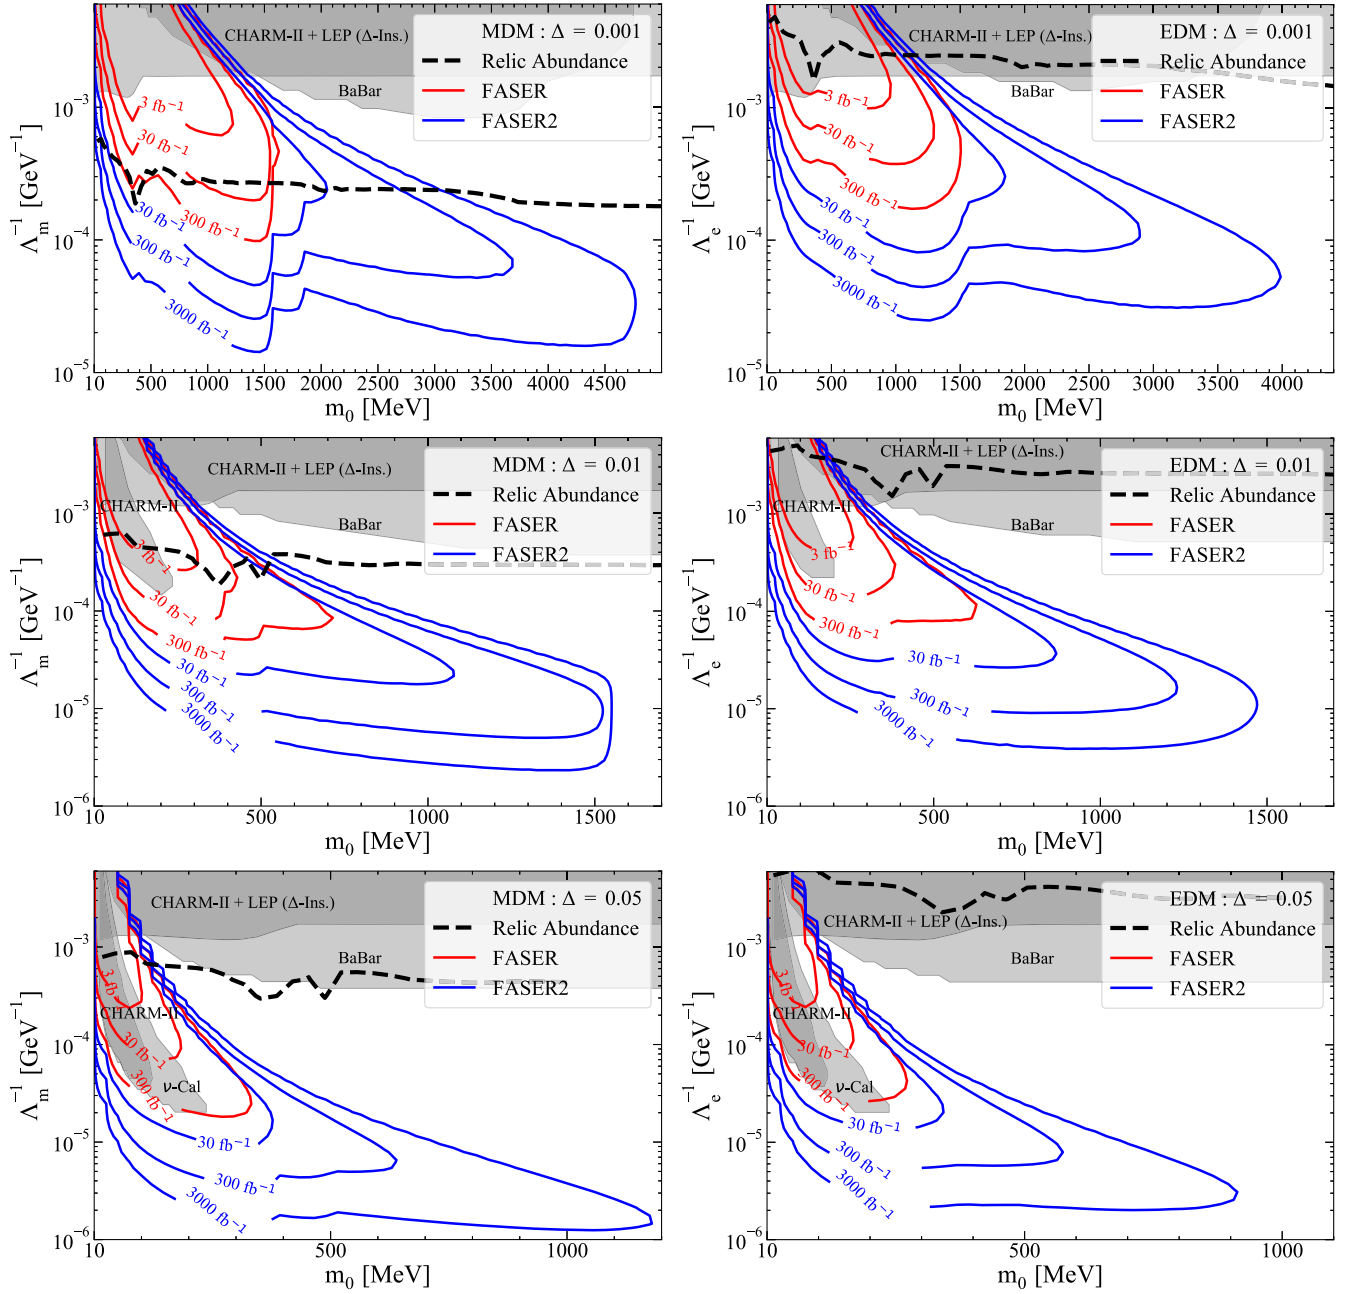


FIG. 3. The integrated luminosity required to observe a monophoton signal of inelastic-dipole dark matter at 95% C.L. at FASER (red curves) and FASER2 (blue curves) within the (Λ_0^{-1}, m_0) -plane for a variety of different choices of Δ . The top, middle, and bottom panels in each column respectively correspond to the choices $\Delta = 0.001$, $\Delta = 0.01$, and $\Delta = 0.05$, while the left and right columns respectively correspond to the cases in which the dark-sector particles couple to the visible sector via an MDM interaction and via an EDM interaction. The dashed curve in each panel indicates the contour along which the present-day relic abundance of χ_0 accords with the observed dark-matter abundance. The gray regions in each panel are excluded by existing constraints. These include constraints from searches at *BABAR*, constraints from Nu-Cal and CHARM-II limits on χ_1 decay, and constraints from probes of dipole-interacting LLPs that are insensitive to the value of Δ and thus effectively identical to those obtained for an elastic-dipole interaction [30]. The exclusion region labeled “CHARM-II + LEP (Δ -Ins.)” in each panel is excluded by this last set of constraints.

particles reach the FASER decay volume before they decay to take the form $\Lambda_0^{-1} \propto m_0^{-2}$. Indeed, the shape of the upper part of each sensitivity contour shown in each panel of Fig. 3 takes this form.

By contrast, the upper bound on the corresponding range of Λ_0 values for which FASER or FASER2 is sensitive for a given integrated luminosity and a given combination of m_0 and Δ stems from the fact that for very large Λ_0 ,

most χ_1 particles that enter the decay volume of the detector pass through it without decaying. For $\bar{d} \gg L$, the expression for the decay probability in Eq. (3.6) reduces to $P(\bar{d}, h, L) \approx h/\bar{d}$. Moreover, the production rate of χ_1 , which is dominated by the contribution from vector-meson decay, is also suppressed for large $\Lambda_{\mathcal{O}}$ by a factor of $\text{BR}(V \rightarrow \chi_1 \bar{\chi}_0) \propto \Lambda_{\mathcal{O}}^{-2}$. As a result, the expected number of χ_1 decays within the FASER or FASER2 detector scales with $\Lambda_{\mathcal{O}}$, m_0 , and Δ like

$$N_{\chi_1} \propto \text{BR}(V \rightarrow \chi_1 \bar{\chi}_0) P(\bar{d}, h, L) \sim \frac{m_0^4 \Delta^3}{\Lambda_{\mathcal{O}}^4}. \quad (3.8)$$

Thus, for a fixed value of Δ , the maximum value of $\Lambda_{\mathcal{O}}$ for which N_{χ_1} drops below the sensitivity threshold scales with m_0 according to a proportionality relation of the form $\Lambda_{\mathcal{O}} \propto m_0$. This scaling behavior accounts for the shape of the lower part of each sensitivity contour shown in the figure.

The ‘‘bumps’’ that appear in these sensitivity contours—bumps which are the most prominent in the upper left panel—arise as a consequence of the decay kinematics of the individual meson species that contribute to the production of χ_1 . In particular, as m_0 increases, mesons with masses $m_M < (2 + \Delta)m_0$ become kinematically forbidden from decaying into χ_i pairs, and as a result the overall production rate of χ_1 particles decreases.

Comparing the results shown in the corresponding panels in the left and right columns of the figure, we observe that the discovery reach obtained for an MDM operator structure at both FASER and FASER2 is comparable to, but nevertheless slightly larger than, the reach obtained for an EDM operator structure. This is a reflection of the fact that as m_0 approaches the kinematic cutoff $m_0 = m_V/(2 + \Delta)$ above which the decay process $V \rightarrow \chi_1 \bar{\chi}_0$ is kinematically forbidden for a given vector-meson species, $\sigma_{\chi_1 M}^{(\text{eff})}$ falls off more rapidly for an EDM operator structure than it does for a MDM operator structure, as indicated in Fig. 1.

The primary message of Fig. 3, however, is that both FASER and FASER2 are capable of probing a sizable region of the parameter space of inelastic-dipole-dark-matter models that has not been meaningfully probed by other experiments. As one would expect, the discovery reach afforded by FASER2 within that parameter space is significantly larger than that afforded by FASER. However, we observe that both experiments are sensitive within hitherto unprobed regions of that parameter space wherein the χ_0 particle has the correct relic abundance to account for the entirety of the dark matter. Moreover, we emphasize that while regions of parameter space that lie below the black dashed line would overproduce dark matter in the context of the standard cosmology, such regions of parameter space would be viable in modified cosmological scenarios involving, for example, additional entropy

production from the out-of-equilibrium decays of heavy particles after thermal freeze-out, but before the nucleosynthesis epoch.

IV. RELIC ABUNDANCE

In order for the χ_i to play the role of the dark matter in our inelastic-dipole scenario, the collective present-day abundance of these states must be consistent with the dark-matter abundance $\Omega_{\text{DM}} \approx 0.26$ inferred from Planck data [32]. Thermal freeze-out provides a natural mechanism for generating abundances for both χ_0 and χ_1 in this scenario. Within the region of parameter space relevant for FASER, χ_1 is sufficiently short-lived that essentially all χ_1 particles present in the universe at the end of the freeze-out epoch decay to χ_0 particles well before present time—and for that matter, well before the big-bang nucleosynthesis epoch. Thus, within this region of parameter space, only χ_0 contributes to the present-day abundance of dark matter. Nevertheless, as we shall see, co-annihilation processes involving χ_1 particles during the freeze-out epoch can have a significant impact on the resulting χ_0 abundance at late times.

The evolution of the number densities n_i of the χ_i in the early universe is governed by a set of coupled Boltzmann equations, each of which takes the form

$$\frac{dn_i}{dt} + 3Hn_i = C_i, \quad (4.1)$$

where H is the Hubble parameter and where C_i represents the contribution to the rate of change of n_i from scattering and decay processes. Equivalently, these differential equations may be formulated in terms of the comoving number densities $Y_i \equiv n_i/s$ of the dark-sector species, where s denotes the entropy density of the universe. Indeed, during periods wherein the total entropy $S \propto sa^3$ within a comoving volume is conserved, the Boltzmann equations for the Y_i each take the particularly simple form

$$\frac{dY_i}{dt} = \frac{C_i}{s}. \quad (4.2)$$

In inelastic-dipole-dark-matter models, the processes that can contribute significantly to the C_i include the following, where $\ell = e, \mu, \tau$ denotes a charged-lepton species:

- (i) t -channel annihilation processes $\bar{\chi}_i \chi_i \leftrightarrow \gamma\gamma$;
- (ii) s -channel co-annihilation processes $\bar{\chi}_i \chi_j \leftrightarrow \ell^+ \ell^-$ and $\bar{\chi}_i \chi_j \leftrightarrow \text{hadrons}$, where $i \neq j$;
- (iii) up- and down-scattering processes $\chi_0 \ell^\pm \leftrightarrow \chi_1 \ell^\pm$ and analogous processes involving hadrons;
- (iv) decay and inverse-decay processes $\chi_1 \leftrightarrow \chi_0 \gamma$.

The annihilation and co-annihilation processes listed above can change the *total* comoving number density $Y_{\text{tot}} \equiv Y_0 + Y_1$ of dark-sector particles. By contrast,

the other processes listed above can transfer comoving number density from one dark-sector species to the other, but have no effect on Y_{tot} .

We can express the C_i for each χ_i as a sum of four terms, each associated with one of the four classes of process itemized above. In particular, C_0 and C_1 may be written in the form

$$\begin{aligned} C_0 &= -[n_0^2 - (n_0^{\text{eq}})^2] \langle \sigma_{00} v \rangle - (n_0 n_1 - n_0^{\text{eq}} n_1^{\text{eq}}) \langle \sigma_{01} v \rangle \\ &\quad - \left(n_0 - \frac{n_1}{n_1^{\text{eq}}} n_0^{\text{eq}} \right) n_e^{\text{eq}} \langle \sigma_{0e} v \rangle + \left(n_1 - \frac{n_0}{n_0^{\text{eq}}} n_1^{\text{eq}} \right) \langle \Gamma \rangle, \\ C_1 &= -[n_1^2 - (n_1^{\text{eq}})^2] \langle \sigma_{11} v \rangle - (n_0 n_1 - n_0^{\text{eq}} n_1^{\text{eq}}) \langle \sigma_{01} v \rangle \\ &\quad + \left(n_0 - \frac{n_1}{n_1^{\text{eq}}} n_0^{\text{eq}} \right) n_e^{\text{eq}} \langle \sigma_{0e} v \rangle - \left(n_1 - \frac{n_0}{n_0^{\text{eq}}} n_1^{\text{eq}} \right) \langle \Gamma \rangle, \end{aligned} \quad (4.3)$$

where n_i^{eq} is the equilibrium number density of χ_i , where n_ℓ^{eq} is the equilibrium number density of the charged-lepton species ℓ , and where $\langle \Gamma \rangle$ denotes the thermally averaged decay rate for χ_1 . In addition, $\langle \sigma_{00} v \rangle$, $\langle \sigma_{11} v \rangle$, $\langle \sigma_{12} v \rangle$, and $\langle \sigma_{0\ell} v \rangle$ denote the thermally averaged cross-sections for the annihilation of two χ_0 particles, the annihilation of two χ_1 particles, the co-annihilation of a χ_0 and a χ_1 particle, and the up-scattering process $\chi_0 \ell^\pm \rightarrow \chi_1 \ell^\pm$, respectively. Expressions for the thermally averaged cross-sections and decay widths for all relevant processes are provided in Appendix B.

Within our parameter-space region of interest, χ_0 and χ_1 always depart from thermal equilibrium only after they have both become nonrelativistic. Within this region of interest, whether or not co-annihilation processes have a significant impact on the overall dark-matter abundance depends primarily on the value of Δ . When $\Delta \gtrsim 1$, co-annihilation processes are relatively unimportant because the number density of χ_1 will be exponentially suppressed relative to that of χ_0 at the time when χ_1 effectively becomes nonrelativistic. As a result, the cosmological evolution of χ_1 is essentially decoupled from that of χ_0 . By contrast, when $\Delta \ll 1$, the number densities of χ_0 and χ_1 remain comparable until much later, and the effect of co-annihilation processes cannot be ignored. Indeed, within the region of parameter space relevant for FASER, wherein $\Delta \ll 1$, co-annihilation processes turn out to play a significant role—and indeed often a dominant role—in determining the overall present-day abundance of χ_0 .

Of course, the freeze-out dynamics in this scenario, and therefore the late-time abundance of χ_0 , depends not only on Δm , but on m_0 as well. Since the temperature at which thermal freeze-out occurs is typically $m_0/\mathcal{O}(10)$, for small masses $m_i \ll m_\mu$ the production of χ_i involves only electrons and positrons. However, for $m_i \gtrsim m_\mu$, processes involving muons are also relevant. At even

higher temperatures, interactions involving hadrons or tau leptons become relevant as well.

At first glance, it might seem that all of the collision terms in Eq. (4.3) contribute to maintaining thermal equilibrium between the χ_i and the thermal bath of SM particles, given that all of these terms vanish when $n_i = n_i^{\text{eq}}$. Were this the case, the up- and down-scattering process in which the χ_i participate, as well as the decay and inverse-decay processes, would serve to keep Y_0 and Y_1 in equilibrium for a significant duration, given that the number densities of the light SM particles that participate in these processes are still quite large during the freeze-out epoch. As a result, the relic abundance of χ_0 would be heavily suppressed. However, as we have discussed above, these processes—unlike annihilation and co-annihilation processes—have no effect on Y_{tot} . Thus, after the annihilation and co-annihilation rates drop below H , both χ_0 and χ_1 depart from the thermal equilibrium with the SM thermal bath. Nevertheless, the up- and down-scattering processes, together with the decay and inverse-decay processes, serve to maintain the relations

$$\begin{aligned} \frac{Y_0}{Y_1} &= \frac{Y_0^{\text{eq}}}{Y_1^{\text{eq}}}, \\ Y_{\text{tot}} &\equiv Y_0 + Y_1 \approx \text{const.} \end{aligned} \quad (4.4)$$

between Y_0 and Y_1 even after this departure from equilibrium occurs. These relations imply that

$$\begin{aligned} Y_0 &= \frac{Y_0^{\text{eq}}}{Y_0^{\text{eq}} + Y_1^{\text{eq}}} Y_{\text{tot}}, \\ Y_1 &= \frac{Y_1^{\text{eq}}}{Y_1^{\text{eq}} + Y_0^{\text{eq}}} Y_{\text{tot}}. \end{aligned} \quad (4.5)$$

Thus, at late time, when $T \ll m_1 - m_0$ and the equilibrium abundances therefore satisfy $Y_1^{\text{eq}} \ll Y_0^{\text{eq}}$, it therefore follows that

$$\begin{aligned} Y_0 &\approx Y_{\text{tot}}, \\ Y_1 &\approx \left(\frac{m_1}{m_0} \right)^{3/2} \exp\left(-\frac{m_1 - m_0}{T} \right) Y_{\text{tot}}. \end{aligned} \quad (4.6)$$

Physically, the results in Eq. (4.6) reflect the fact that even if Y_0 and Y_1 are comparable at the moment when χ_0 and χ_1 initially depart from thermal equilibrium with the radiation bath, up-scattering and inverse-decay processes will become increasingly kinematically disfavored as T drops. As a result, the initial population of χ_1 particles present at the point at which this departure occurs will gradually be converted into χ_0 particles via decay and down-scattering processes until the majority of this initial population is all but depleted and χ_0 freezes out.




FIG. 4. The evolution of the comoving number densities Y_i of the two dark-sector fields χ_i , shown as functions of m_0/T for the parameter assignments specified in the legend. The solid red and dashed blue curves respectively represent Y_0 and Y_1 , whereas the solid orange and dashed cyan curves represent the corresponding equilibrium comoving number densities Y_0^{eq} and Y_1^{eq} . The dotted gray line indicates the value of Y_i that yields the correct present-day dark-matter abundance for this choice of model parameters. The yellow curve represents the quantity $Y_0(Y_1^{\text{eq}}/Y_0^{\text{eq}})$, which, in accord with Eq. (4.4), effectively coincides with Y_1 across the entire range of m_0/T shown.

In Fig. 4, we show how the comoving number densities Y_0 and Y_1 evolve as a function of m_0/T for an example choice of model parameters that yields a relic abundance for χ_0 that accords with Ω_{DM} . Since $\Delta \ll 1$ for this choice of parameters, the equilibrium comoving number densities Y_0^{eq} and Y_1^{eq} (indicated by the solid orange and dashed cyan curves in the figure) are very similar. As a result, the actual comoving number densities Y_0 and Y_1 , which track their equilibrium values very closely at high temperatures, likewise remain quite similar at high temperatures. However, at around $m_0/T \lesssim 20$, when Y_0 and Y_1 begin to depart appreciably from their equilibrium values, they also begin to differ appreciably from each other. Indeed, by the time at which Y_0 has effectively settled into its late-time asymptotic value—a time that corresponds roughly to the point at which $m_0/T \sim 200$ —this comoving number density is already an order of magnitude larger than Y_1 . We also observe that Y_1 tracks the quantity $Y_0(Y_1^{\text{eq}}/Y_0^{\text{eq}})$, in accord with the result in Eq. (4.4).

We also note that the value of $\Lambda_{\mathcal{O}}$ required to obtain the correct thermal relic density in the EDM case is smaller than it is in the MDM case for the same choice of m_0 and Δ . This is ultimately a consequence of angular-momentum conservation. In the regime in which $\Delta \ll 1$ and χ_0 and χ_1 are nearly degenerate in mass, the thermal relic density is determined primarily by co-annihilation processes such as $\bar{\chi}_0\chi_1 \rightarrow e^+e^-$, which are mediated by an s -channel photon. In the MDM case, these are S -wave processes: as noted in

Sec. III A, the initial state has spin $S = 1$, and since the final state also has spin $S = 1$, angular-momentum conservation can be satisfied by states with orbital angular momentum $L = 0$. By contrast, in the EDM case, they are P -wave processes: the initial state has spin $S = 0$ and the final state has spin $S = 1$, so angular-momentum conservation requires orbital angular momentum $L = 1$. The value of $\Lambda_{\mathcal{O}}$ must therefore be smaller in the EDM case to compensate for the resulting P -wave suppression.

We can obtain a more quantitative picture of how this P -wave suppression affects the freeze-out dynamics in the EDM case relative to the MDM case by comparing the corresponding expressions for the cross-section $\sigma_{01 \rightarrow ee}^{\text{MDM}}$ and $\sigma_{01 \rightarrow ee}^{\text{EDM}}$ for the co-annihilation process $\bar{\chi}_0\chi_1 \rightarrow e^+e^-$ —expressions respectively obtained by substituting the $C^{(s)}$ factors in Eqs. (B9) and (B9) into Eq. (B8). In particular, since $\Delta \ll 1$ within our parameter-space region of interest, one can gain some insight into how the differences between these expressions affect our results by considering how these expressions behave in the $\Delta \rightarrow 0$ limit. In this limit, one finds that these expressions reduce to

$$\begin{aligned} \sigma_{01 \rightarrow ee}^{\text{MDM}} &= \frac{e^2}{6\pi\Lambda_m^2} \left(\frac{3 - 2v^2}{v} \right), \\ \sigma_{01 \rightarrow ee}^{\text{EDM}} &= \frac{e^2}{6\pi\Lambda_e^2} v, \end{aligned} \quad (4.7)$$

where $v = \sqrt{1 - 4m_0^2/s}$ is the speed of either initial-state particle in the center-of-mass frame. Thus, we see that the co-annihilation cross-section in the EDM case is suppressed relative to the co-annihilation cross-section in the MDM case by a factor of v^2 . At freeze-out, when $v^2 \sim 0.1$, the two cross-sections are equal when $\Lambda_m \sim 5\Lambda_e$. Indeed, this is reflected in the locations of the $\Omega_{\chi_0} = \Omega_{\text{DM}}$ contours appearing in Fig. 3.

V. EXISTING CONSTRAINTS

The parameter space of inelastic-dipole-dark-matter models is constrained by a variety of experimental and observational considerations. Within the region of that parameter space that we have identified as being the most relevant for FASER, the most important such constraints are those from colliders and from fixed-target experiments. In addition, in cases in which the present-day abundance of χ_0 particles is non-negligible, astrophysical considerations also imply constraints on that region of parameter space. We discuss these constraints in turn below.

A. Decay signatures: Colliders

The *BABAR* experiment at the PEP-II e^+e^- facility is sensitive to regions of parameter space near the thermal-relic contour for inelastic-dipole dark matter and indeed can be seen as complementary to FASER. The resulting

constraints on this parameter space were investigated in Ref. [25] for the case of a MDM interaction. We extend this analysis here both to the case of an EDM interaction and to smaller values of Δ .

Limits on events involving either one or two energetic photons and sizable missing energy \cancel{E} at *BABAR* place constraints on inelastic-dipole dark matter. The dark-sector particles χ_0 and χ_1 can be produced there either through the t -channel process $e^+e^- \rightarrow \gamma\chi_1\bar{\chi}_0$ or through the s -channel process $e^+e^- \rightarrow \chi_1\bar{\chi}_0$. The leading contribution to the signal-event rate in the $2\gamma + \cancel{E}$ channel arises from t -channel production, followed by the prompt decay of the χ_1 . By contrast, sizable contributions to the signal-event rate in the $\gamma + \cancel{E}$ channel arise both from s -channel production, followed by the prompt decay of the χ_1 , and from t -channel production, followed by the χ_1 escaping the detector before it decays.

For our analysis of the signal rate in both the $2\gamma + \cancel{E}$ and $\gamma + \cancel{E}$ channels, we make use of the *BABAR* monophoton trigger, which was in effect while 60 fb^{-1} of integrated luminosity was accumulated. This trigger requires that the leading photon in the event have energy $E_\gamma \geq 2 \text{ GeV}$. For both channels, we also require that $\cancel{E} \geq 20 \text{ MeV}$. For the $2\gamma + \cancel{E}$ channel, following Refs. [25,29], we also require that an additional photon with $E_\gamma \geq 50 \text{ MeV}$ be present in the event.

For the $2\gamma + \cancel{E}$ channel, we assume no background and require that the χ_1 decays promptly (i.e., within 1 cm of the interaction point). Since any χ_1 particles which might have been detected at *BABAR* would necessarily have been produced with much smaller boosts, this prompt decay length is the one relevant for our parameter-space region of interest. By contrast, sizable SM backgrounds to the $\gamma + \cancel{E}$ signature arise from events involving neutrino production via neutral-current interactions and from events in which a photon is produced in conjunction with additional photons or e^+e^- pairs that escape detection. We find that the constraints derived from the $\gamma + \cancel{E}$ channel turn out to be subdominant to the constraints from the $2\gamma + \cancel{E}$ channel, and thus we focus on these constraints in what follows.

The signal contribution in the $2\gamma + \cancel{E}$ channel involves the prompt decay of the χ_1 , and so the corresponding constraints on our model-parameter space extend to arbitrarily small $\Lambda_{\mathcal{O}}$. These constraints exclude the shaded region of model-parameter space labeled “*BABAR*” in each panel of Fig. 3. We note that the Belle-II experiment [40] at SuperKEKB, which is currently taking data, will be able to extend the reach of B-factory experiments within this parameter space to larger $\Lambda_{\mathcal{O}}$ in the near future.

Since χ_1 can also be produced via analogous processes at hadron colliders, LHC searches likewise place constraints on the parameter space of inelastic-dipole-dark-matter models. The extent to which such searches are capable of probing the parameter space of such models was investigated in Ref. [25] for the case of a MDM interaction.

While the results of this analysis indicate that LHC searches are in general capable of probing that parameter space down to $\Lambda_{\mathcal{O}} \sim \mathcal{O}(10^3) \text{ GeV}$ for a broad range of m_0 and Δ , they also indicate that for $\Delta \lesssim 0.05$, the region of parameter space that these searches are capable of probing does not extend beyond the region that is already excluded by *BABAR* data. Thus, the constraints from LHC searches on the region of that parameter space relevant for FASER are subdominant to those from *BABAR* and need not be considered further.

One could also consider constraints from the LEP experiment constraining invisible Z decays. Such constraints were considered in Ref. [25], where the authors explored a similar inelastic dipole operator, but with a coupling to the hypercharge boson instead of the photon. For the magnetic dipole case, they derived a lower bound $\Lambda_m \gtrsim 10^3 \text{ GeV}$ on the effective scale associated with the corresponding operator. Depending on the UV theory, a similar constraint might be applicable to our model as well. However since the χ_i have no direct coupling to the physical Z boson in the effective theory we consider in this paper, we do not consider this constraint further.

B. Decay signatures: Fixed-target experiments

A variety of experiments in which a beam of SM particles—typically electrons or protons—collides with a fixed target located some distance in front of a particle detector impose nontrivial constraints on inelastic-dipole-dark-matter models. Indeed, χ_1 particles can be produced by these collisions, and some fraction of these χ_1 particles subsequently decay within the detector, producing photons that can then be detected. At electron beam-dump experiments, χ_1 particles are produced predominately via bremsstrahlung-like processes such as $e^-N \rightarrow e^-N\chi_1\bar{\chi}_0$, where N denotes an atomic nucleus. At proton beam dumps, these particles are produced primarily via the decay of various mesons produced by the collisions between the incident protons and the nuclei in the target. We shall consider the constraints from these two classes of fixed-target experiments in turn.

Among electron beam-dump experiments, E137 [41] and E141 [42] are the most relevant for constraining the region of model-parameter space pertinent for FASER. At each of these experiments, electrons with lab-frame energies $E_e \sim \mathcal{O}(10) \text{ GeV}$ collided with a dense target—aluminum in the case of E137, tungsten in the case of E141. These experiments also had similar, $\mathcal{O}(\text{GeV})$ detection thresholds. By contrast, the beam-dump experiment at Orsay [43], which had a similar detection threshold but a far lower electron-beam energy $E_e = 1.6 \text{ GeV}$, yields far weaker constraints within this same region of parameter space. The beam-dump experiment at KEK [44] had an electron-beam energy $E_e = 2.5 \text{ GeV}$, which is also much lower than that of either E137 or E141. However, this experiment also had a much lower detection

threshold—around $\mathcal{O}(100 \text{ MeV})$. Nevertheless, given the significant background from photon bremsstrahlung at energies $E_\gamma \lesssim 1 \text{ GeV}$, this lower detection threshold is unlikely to result in a significant improvement in sensitivity to monophoton signals. Moreover, the total number of electron collisions attained over the lifetime of the KEK experiment is a factor of 10^3 smaller than the number attained at E137. As a result, we do not expect the KEK experiment to yield constraints competitive with those from E137 within our parameter-space region of interest.

We assess the sensitivity of E137 and E141 within the parameter space of our inelastic-dark-matter model in a manner analogous to that employed in similar studies of other LLP scenarios in Refs. [42,45,46]. In particular, we simulate the production of photons at these experiments using the MG5_AMC@NLO code package [35] and incorporate the appropriate form factors for the target material in each experiment—form factors that account for the effects of coherent scattering as well as nuclear and atomic screening [46–48]. We find that neither E137 nor E141 is capable of achieving a statistically significant number of events necessary to constrain the region of model-parameter space relevant for FASER, even for zero background. This is primarily due to the fact that kinematic considerations place an upper bound $E_\gamma \leq [E_e + (E_e^2 - m_0^2)^{1/2}]\Delta$ on the lab-frame energy of a photon produced by $\chi_1 \rightarrow \chi_0\gamma$ decay at fixed-target experiments of this sort. Since $\Delta \ll 0.1$ within our parameter-space region of interest, E_γ typically lies significantly below the threshold for detection. Indeed, it is only for $\Delta > 0.1$ that the E137 and E141 data become constraining for inelastic-dipole dark matter.

Among proton beam-dump experiments, Nu-Cal [49,50] and CHARM-II [51,52] are the most relevant within our parameter-space region of interest. The Nu-Cal experiment, which primarily functioned as a neutrino detector, collided approximately 10^{18} protons—each with a lab-frame energy of $E_p \approx 69 \text{ GeV}$ —with an iron target over the course of its operation. Scalar and vector mesons produced in these collisions could decay into χ_1 particles via the processes discussed in Sec. III A. Decays of these χ_1 particles to photons within the Nu-Cal detector could then give rise to a signal.

In assessing the expected number of signal events at Nu-Cal in inelastic-dipole-dark-matter models, we adopt the procedure for modeling the scalar- and vector-meson production cross-sections used in Ref. [50]. We relate the inclusive cross-section $\sigma(p\text{Fe} \rightarrow \pi^0 X)$ for π^0 production via proton scattering with an iron nucleus N to the inclusive cross-section for π^0 production via proton-proton scattering

$$\sigma(pN \rightarrow \pi^0 X) = A^\alpha \sigma(pp \rightarrow \pi^0 X), \quad (5.1)$$

where X denotes any additional SM particles in the final state, where $A = 56$ is the atomic mass number of iron, and

where $\alpha = 0.55$. We also model the differential cross-section for π^0 production via proton-proton scattering as an average of the corresponding differential cross-sections for π^+ and π^- production, which have been measured empirically [53]. In other words, we take

$$\frac{d^2\sigma(pp \rightarrow \pi^0 X)}{dzd(p_T^2)} = \frac{1}{2} \left[\frac{d^2\sigma(pp \rightarrow \pi^- X)}{dzd(p_T^2)} + \frac{d^2\sigma(pp \rightarrow \pi^+ X)}{dzd(p_T^2)} \right], \quad (5.2)$$

where z is the fraction of the outgoing momentum of the incident proton carried by the π^0 and where p_T is the magnitude of the component of its momentum vector transverse to the beam axis. Finally, following Ref. [31], we estimate the differential cross-sections for the production of all heavier neutral scalar- and vector-meson species M as

$$\frac{d^2\sigma(pN \rightarrow MX)}{dzd(p_T^2)} = \mathcal{N}_M^{(\text{POT})} \frac{d^2\sigma(pN \rightarrow \pi^0 X)}{dzd(p_T^2)}, \quad (5.3)$$

where the scaling factor $\mathcal{N}_M^{(\text{POT})}$ represents the average number of mesons of that species produced per proton on target. Conservative estimates for the $\mathcal{N}_M^{(\text{POT})}$ values for all relevant meson species were computed in Ref. [31] using PYTHIA 8.2 [54,55] simulations, and we adopt these $\mathcal{N}_M^{(\text{POT})}$ values as well.

Given these differential cross-sections, we may calculate the spectrum of χ_1 particles produced in the direction of the Nu-Cal detector by the decays of these mesons in a straightforward manner. This detector consisted of a cylindrical decay volume 23 m in length and 2.3 m in diameter located 64 m behind the target with its axis aligned with the axis of the proton beam. In assessing the total expected number of signal events from $\chi_1 \rightarrow \chi_0\gamma$ decays at Nu-Cal for any combination of $\Lambda_\mathcal{O}$, m_0 , and Δ , we require not only that the χ_1 decay within this decay volume, but also that the lab-frame energy of the resulting photon satisfy $E_\gamma > 3 \text{ GeV}$ —a criterion that derives from the fact that Nu-Cal was only capable of distinguishing an electromagnetic from a hadronic shower when the total energy associated with the shower exceeds this threshold.

By comparing the expected number of signal events for different combinations of $\Lambda_\mathcal{O}$, m_0 , and Δ to the number of events satisfying the same criteria which were actually observed at Nu-Cal during its full run—a total of 5 such events were observed, and the expected background was 3.5 events—we can derive constraints on our model-parameter space. These constraints exclude the shaded region of model-parameter space labeled “ ν -Cal” in the bottom two panels of Fig. 3, which correspond to the parameter choice $\Delta = 0.05$. Similar exclusion regions do not appear in the other panels of the figure, however, because too few photons satisfy the $E_\gamma > 3 \text{ GeV}$ criterion

for such small values of Δ . Since the vector mesons ρ and ω , which have the highest production rates at Nu-Cal, are kinematically forbidden from decaying to χ_i pairs for $m_0 \gtrsim 390$ MeV, the results of this experiment have little constraining power within our parameter space for m_0 above this value.

The CHARM experiment and its successor CHARM-II [51,52] would likewise have been able to produce χ_i pairs through meson decay. Thus, the results of these experiments also constrain the parameter space of inelastic-dipole-dark-matter models. The original CHARM experiment collided approximately 7.1×10^{18} protons—each with an energy of $E_p = 450$ GeV—with a copper target over the course of its operation. The center of the detector, which was situated 487.3 m behind the target, had a transverse area of 2.4×2.4 m². Following [56], which searched for a monophoton signal, we take the fiducial decay volume to be $2.4 \times 2.4 \times 11.2$ m³. The CHARM II experiment collided approximately 2.5×10^{19} protons—each also with an energy of $E_p = 450$ GeV—with a beryllium target over the course of its operation. The detector, which was situated 870 m behind the target, had a transverse area of 3.7×3.7 m². We take the length of the fiducial decay volume to be 35.7 m, which represents the length of the entire calorimeter [57], in the direction parallel to the beam axis. While the distance between the target and the detector at each of these experiments was similar to the distance between FASER and the ATLAS IP, the characteristic lab-frame energy of a χ_1 particle produced by collisions of the proton beams at CHARM and CHARM-II with their respective fixed targets was significantly lower than the characteristic lab-frame energy of a χ_1 particle produced in the forward direction at the LHC. As a result, CHARM and CHARM-II data constrain larger values of $\Lambda_{\mathcal{O}}$ than does FASER data, which lead to shorter proper lifetimes for χ_1 .

In deriving the constraints from CHARM data on our parameter space, we largely follow the procedure outlined in Refs. [31,56]. Likewise, in assessing the corresponding constraints from CHARM-II data, we largely follow the procedure outlined in Ref. [31]. For each experiment, we derive a differential cross-section for π^0 production of the BMPT form [58] using the BDNMC code package [59]. We then obtain the differential cross-sections for the production of η , η' , ρ , ω , ϕ , and J/ψ by scaling this differential cross-section for π^0 production by an overall factor $\mathcal{N}_M^{(\text{POT})}$, as in Eq. (5.3).

From these differential cross-sections, we may derive an expected number of signal events from χ_1 decays within the decay volume at either CHARM or CHARM-II for a given combination of $\Lambda_{\mathcal{O}}$, m_0 , and Δ in much the same way as we derived the corresponding signal-event count at Nu-Cal. In accord with the cuts performed on the energy of the electromagnetic shower in Ref. [60], we require that the

lab-frame energy of the photon satisfy $7.5 \text{ GeV} \leq E_\gamma \leq 50 \text{ GeV}$ for the CHARM analysis. Likewise, for the CHARM-II analysis, we require that $3 \text{ GeV} \leq E_\gamma \leq 24 \text{ GeV}$ in accord with the corresponding cuts performed in Refs. [51,57]. Although the minimum value of E_γ for both of these windows is considerably higher than the threshold value employed in the Nu-Cal analysis above, E_p is also considerably higher for both CHARM and CHARM-II than it is for Nu-Cal. As a result, the characteristic energies of the χ_1 particles are also significantly higher, and a non-negligible fraction of photons produced by $\chi_1 \rightarrow \chi_0 \gamma$ decay survive these E_γ cuts, even for $\Delta = 0.01$.

We derive constraints on our model-parameter space by comparing the expected number of signal events obtained for different combinations of $\Lambda_{\mathcal{O}}$, m_0 , and Δ to the number of events that were actually observed at CHARM or CHARM-II. However, there are significant uncertainties in the expected backgrounds at these experiments [51,56,60]. Thus, since our primary aim in this paper is to identify regions of our parameter space within which FASER and FASER2 are sensitive, but which are clearly *not* constrained by other experiments, we maximize the size of the exclusion regions from CHARM and CHARM-II by assuming a background of zero events at each experiment. The resulting constraint from CHARM-II data excludes the shaded region of model-parameter space labeled “CHARM-II” in each panel of Fig. 3. Since we find that these regions completely subsume those excluded by CHARM data, we do not include separate exclusion contours for CHARM in this figure.

C. Constraints from the elastic limit

In addition to the signals discussed above, which arise as a consequence of χ_1 decay and therefore vanish in the elastic limit, i.e., the limit in which $\Delta \rightarrow 0$, there are a number of potential signatures of dipole-interacting dark matter at colliders, fixed-target experiments, and neutrino facilities that are associated with processes that do not vanish in that limit. These signatures include, for example, electron or nucleon recoils precipitated by collisions of the χ_i with the detector material in fixed-target experiments and monojet + \cancel{E} signatures at hadron colliders. The nonobservation of such signatures at experiments including E613 [61], MiniBooNE [62], LSND [63], LEP [64], CHARM-II [51], and LSND [63] places nontrivial constraints on inelastic-dipole-dark-matter models.

Since the mass-splittings $m_1 - m_0$ within our parameter-space region of interest are quite small in comparison with the characteristic energy scales of these experiments, the corresponding constraints are well approximated by those obtained for a purely elastic dipole interaction. Detailed studies of the experimental constraints on elastic dipole-interacting dark matter have been performed for both the MDM and EDM cases [29–31]. Throughout most of our parameter-space region of interest, the leading

such constraints are those from limits on electron recoils at CHARM-II or from limits on monophoton + \cancel{E} searches at LEP. The shaded region of model-parameter space labeled “CHARM-II + LEP (Δ -Ins)” in each panel of Fig. 3 is excluded by constraints of this sort. These constraints collectively impose a bound $\Lambda_{\mathcal{O}} \gtrsim 1$ TeV within our parameter-space region of interest that is not only independent of Δ , but also not particularly sensitive to the value of m_0 .

D. Astrophysical constraints

In cases in which χ_0 particles have a significant present-day abundance, their annihilation within the halos of galaxies via the t -channel process $\chi_0\bar{\chi}_0 \rightarrow \gamma\gamma$ can give rise to a monochromatic photon signal at gamma-ray telescopes. The leading constraints on such a signal are currently those from Fermi-LAT data, which imply an upper bound on the corresponding thermally averaged annihilation cross-section that varies with m_0 from $\langle\sigma v\rangle \lesssim 10^{-34}$ cm³ s⁻¹ for $m_0 \approx 1$ MeV to $\langle\sigma v\rangle \lesssim 10^{-29}$ cm³ s⁻¹ for $m_0 \approx 1$ GeV [65]. This thermally averaged cross-section may be estimated as

$$\langle\sigma v\rangle \approx 10^{-33} \text{ cm}^3 \text{ s}^{-1} \left(\frac{10^3 \text{ GeV}}{\Lambda_{\mathcal{O}}}\right)^4 \left(\frac{m_0}{\text{GeV}}\right)^2 \left(\frac{v}{10^{-3}}\right), \quad (5.4)$$

which is well below the Fermi-LAT bound for all m_0 within our parameter-space region of interest. Constraints from other instruments (such as EGRET) on monochromatic gamma-ray signals from the annihilation of sub-GeV dark-matter particles have also been assessed [66], but turn out to be subleading in comparison with the Fermi-LAT bounds.

In principle, an additional constraint on inelastic-dipole-dark-matter models arises from observational limits on distortions of the cosmic microwave background [67] due to annihilation, co-annihilation, or scattering processes involving the χ_i during or after recombination. However, we find that these limits do not constrain any portion of our parameter-space region of interest that is not excluded by other constraints. Indeed, since the cross-section for the annihilation process $\chi_0\bar{\chi}_0 \rightarrow \gamma\gamma$ is proportional to $\Lambda_{\mathcal{O}}^{-4}$, the corresponding annihilation rate is highly suppressed for values $\Lambda_{\mathcal{O}} \gtrsim 1$ TeV that are consistent with constraints from collider and fixed-target experiments. Moreover, the lifetime of the χ_1 is sufficiently short throughout our parameter-space region of interest that the abundance of χ_1 particles is negligible at the time of recombination. At the same time, for $m_0 \gtrsim 1$ MeV and $\Delta \gtrsim 0.001$, the mass-splitting between χ_0 and χ_1 satisfies $m_1 - m_0 \gtrsim 1$ keV. Since this significantly exceeds the temperature $T \sim \mathcal{O}(10 \text{ eV})$ during the recombination epoch, the production of χ_1 by up-scattering processes immediately prior to or during recombination is highly Boltzmann-suppressed.

Taken together, these considerations imply that the co-annihilation rate during recombination is negligible.

Limits on the effective number N_{eff} of neutrino species during the nucleosynthesis epoch place stringent constraints [68] on light thermal relics with highly suppressed couplings to the visible sector. However, such constraints only become relevant for $m_0 \lesssim 10$ MeV, and do not impact the region of model-parameter space relevant for FASER.

Finally, direct-detection experiments also place constraints on the scattering cross-sections of χ_0 with atomic nuclei or with electrons. However, inelastic-scattering rates for a dark-matter particles with $m_0 \sim \mathcal{O}(\text{GeV})$ at such experiments are highly suppressed for $\Delta \gtrsim 10^{-6}$. While elastic scattering via loop-level processes still yield a contribution to the overall scattering rate for larger values of Δ [69], this contribution is far too small within our parameter-space region of interest to be constraining.

VI. CONCLUSIONS AND DIRECTIONS FOR FUTURE RESEARCH

In this paper, we investigated the prospects for detecting a signature of inelastic-dipole dark matter at the FASER detector and its proposed upgrade FASER2. This signature involves the production of the heavier of the two dark-sector states present in the theory by collisions at the ATLAS IP and its subsequent decay into the lighter dark-sector state and a single photon inside the FASER decay volume. We determined the discovery reach of FASER and FASER2 within the parameter space of this model for the case of either a magnetic or an electric dipole-moment interaction and showed that these detectors are capable of probing regions of this parameter space within which other experimental probes have little or no sensitivity. Moreover, these include regions wherein thermal freeze-out yields an abundance for the lighter dark-sector state that agrees with the observed abundance of dark matter.

Several comments are in order. First, it is worth noting that one of the reasons why FASER and FASER2 are capable of probing regions of inelastic-dipole-dark-matter parameter space within which existing experiments have little constraining power is simply the far higher energy of the LHC. Indeed, as we saw in Sec. V, one of the reasons why current fixed-target bounds are not terribly constraining within our parameter-space region of interest is that the majority of photons produced by χ_1 decays within the detector at these experiments would have had lab-frame energies that are below the corresponding detection thresholds. By contrast, the substantial boost provided to χ_1 particles at the LHC ensures that the photons produced by the decays of these particles within the FASER detector have far higher energies. Moreover, the time-dilation factor provided by these large boosts enhances the ability of FASER to probe regions of parameter space in which the lifetime of χ_1 is considerably shorter. Since small mass splittings are a generic possibility that leads to long-lived

particles, and since these small mass splittings in turn lead to soft decay products, one can expect the high energies and boosts available at the LHC to be advantageous for probing many similar LLP models beyond the specific one discussed here. Such models might include, for example, inelastic dark-matter models involving other coupling structures, such as the anapole or charge-radius operators [29–31].

In this connection, it is worth noting that other proposed experiments would also be able to probe parts of our region of interest within the parameter space of inelastic-dipole dark matter models. At CERN, these include the SHiP experiment [27,28], which would collide a 400 GeV proton beam with a fixed target. An estimate of the reach of this experiment within our parameter-space region of interest is provided in Appendix C. As can be seen there, the reach depends critically on the mass splitting between the dark states and the energy threshold for monophoton signal detection. For small Δ , FASER and FASER2 probe regions beyond the reach of SHiP, but for larger Δ , their reaches may be very roughly comparable. Further detailed studies of the energy threshold for monophoton signal detection are clearly needed. In Ref. [13], the authors explored the sensitivity of SHiP and FASER to an inelastic sterile neutrino interacting through an MDM operator and found SHiP to be more constraining than FASER. This is a consequence of their requiring very high thresholds for monophoton detection for FASER ($E_\gamma \geq 10$ GeV), while requiring much lower thresholds at SHiP ($E_\gamma \geq 0.1$ GeV). Although these thresholds are not precisely known for each experiment, FASER has significantly more shielding against background and is likely better suited to finding very low-energy signal photons. We also note that other proposed experiments, such as SHADOWS [70] have the potential to probe this parameter-space region, although the same comments about energy thresholds apply for all fixed target experiments. Other LLP detectors such as ANUBIS [71], CODEX-b [72], MAPP [73], MATHUSLA [74] are primarily designed to be sensitive to LLPs produced with significantly more transverse momenta than those considered here. Thus data from these detectors is not expected to be competitive with data from FASER in the parameter space we consider.

Finally, in this paper we have considered the case of a dark sector that comprises two states χ_0 and χ_1 that interact with each other and with the visible sector via an inelastic-dipole interaction. While this model is interesting in its own right, it can also be viewed as merely the simplest possible realization of a nonminimal dark sector consisting of N different states χ_i with $i = 0, 1, \dots, N-1$ that interact with each other via a set of operators \mathcal{O}_{ij} , each with the MDM or EDM structure given in Eq. (2.1). Within this more general class of models, decay cascades involving multiple sequential LLP decays of the form $\chi_i \rightarrow \chi_j \gamma$ can develop. Scenarios in which decay cascades involving multiple LLPs can arise frequently give rise to distinctive

phenomenological signatures [20] that differ from those observed in more traditional LLP scenarios. It would therefore be interesting to explore the prospects for probing dark-sector scenarios of this sort at FASER, FASER2, and other proposed experiments.

ACKNOWLEDGMENTS

We thank Daniele Barducci, Enrico Bertuzzo, David Casper, Xiaoyong Chu, Patrick deNiverville, Andrei Golutvin, Felix Kling, Dominik Köhler, Jui-Lin Kuo, Gopolang Mohlabeng, Marco Taoso, Claudio Toni, and Sebastian Trojanowski for discussions. The research activities of K. R. D. were supported in part by the U.S. Department of Energy under Grant DE-FG02-13ER41976/DE-SC0009913, and by the U.S. National Science Foundation (NSF) through its employee IR/D program. The work of J. L. F., M. F., and F. H. was supported in part by NSF Grants PHY-1915005 and PHY-2210283. The work of J. L. F. was also supported in part by NSF Grant PHY-2111427, Simons Investigator Award No. 376204, Simons Foundation Grant 623683, and Heising-Simons Foundation Grants No. 2019-1179 and No. 2020-1840. The work of M. F. was also supported in part by NSF Graduate Research Fellowship Award No. DGE-1839285. The work of F.H. was also supported in part by the International Postdoctoral Exchange Fellowship Program. The research activities of S.L. were supported in part by the National Research Foundation of Korea (NRF) grant funded by the Korean government (MEST) (No. NRF-2021R1A2C1005615) and the Samsung Science & Technology Foundation under Project Number SSTF-BA2201-06. The research activities of B. T. were supported in part by the National Science Foundation under Grant PHY-2014104. The opinions and conclusions expressed herein are those of the authors and do not represent any funding agencies.

APPENDIX A: MESON BRANCHING FRACTIONS

In this appendix, we provide explicit formulas for the branching fractions for the meson-decays processes that contribute to χ_1 production in an inelastic-dipole-dark-matter context.

The dominant contribution to χ_1 production from each relevant neutral vector meson V is that associated with the two-body decay process $V \rightarrow \bar{\chi}_0 \chi_1$. It is convenient to parametrize the branching fraction for each such decay process in terms of the corresponding branching fraction for the SM decay process $V \rightarrow e^+ e^-$ as follows:

$$\begin{aligned} \text{BR}(V \rightarrow \bar{\chi}_0 \chi_1) &= \text{BR}(V \rightarrow e^+ e^-) \frac{\Gamma_{V \rightarrow \bar{\chi}_0 \chi_1}}{\Gamma_{V \rightarrow e^+ e^-}} \\ &= \text{BR}(V \rightarrow e^+ e^-) \frac{|\mathcal{M}(V \rightarrow \bar{\chi}_0 \chi_1)|^2 |\vec{\mathbf{p}}_{\chi_0}|}{|\mathcal{M}(V \rightarrow e^+ e^-)|^2 |\vec{\mathbf{p}}_e|}, \end{aligned} \quad (\text{A1})$$

where $|\mathcal{M}(V \rightarrow \bar{\chi}_0\chi_1)|^2$ and $|\mathcal{M}(V \rightarrow e^+e^-)|^2$ are the squared matrix elements for the two decay processes, averaged over initial-state-particle spins and summed over final-state-particle spins, and where $|\vec{\mathbf{p}}_{\chi_0}|$ and $|\vec{\mathbf{p}}_e|$ denote the magnitudes of the three-momenta of either final-state particle in the corresponding process.

The squared matrix element for vector-meson decay to electrons is given by

$$|\mathcal{M}(V \rightarrow e^+e^-)|^2 = \frac{16\pi\alpha}{3} m_V^2 \left(1 + \frac{2m_e^2}{m_V^2}\right). \quad (\text{A2})$$

For the case of an MDM operator structure, the squared matrix element for the process $V \rightarrow \bar{\chi}_0\chi_1$ is given by

$$|\mathcal{M}_{\text{MDM}}(V \rightarrow \bar{\chi}_0\chi_1)|^2 = \frac{8m_V^4}{3\Lambda_m^2} \left[1 + \frac{m_0^2 + 6m_0m_1 + m_1^2}{m_V^2} - \frac{2(m_1^2 - m_0^2)^2}{m_V^4}\right]. \quad (\text{A3})$$

For the case of an EDM operator structure, the squared matrix element for this same process is given by

$$|\mathcal{M}_{\text{EDM}}(V \rightarrow \bar{\chi}_0\chi_1)|^2 = \frac{8m_V^4}{3\Lambda_e^2} \left[1 + \frac{2(m_1 - m_0)^2}{m_V^2}\right] \left[1 - \frac{(m_1 + m_0)^2}{m_V^2}\right]. \quad (\text{A4})$$

We note that for $m_1 = m_0$, the expressions in Eqs. (A3) and (A4) reduce to the corresponding results obtained in Ref. [31] for the case of a purely elastic dipole interaction.

By contrast, the dominant contribution to χ_1 production from each relevant neutral pseudoscalar meson P arises from that associated with the three-body decay process $P \rightarrow \gamma^*\gamma \rightarrow \bar{\chi}_0\chi_1\gamma$, where γ^* denotes an off-shell photon. The coupling between each P and a pair of photons is described by the effective Lagrangian

$$\mathcal{L}_P = \frac{e^2}{16\pi^2 F_P} P F_{\mu\nu} \tilde{F}^{\mu\nu}, \quad (\text{A5})$$

where F_P is the corresponding pseudoscalar decay constant and where $\tilde{F}^{\mu\nu} \equiv \epsilon^{\mu\nu\rho\sigma} F_{\rho\sigma}/2$ is the dual electromagnetic field-strength tensor.

The differential branching fraction with respect to the energy E_1^* of the final-state χ_1 particle in the rest frame of the decaying P can be written in the form

$$\frac{d\text{BR}(P \rightarrow \bar{\chi}_0\chi_1\gamma)}{dE_1^*} = \frac{\text{BR}(P \rightarrow \gamma\gamma)}{\Gamma(P \rightarrow \gamma\gamma)} \times \frac{d\Gamma(P \rightarrow \bar{\chi}_0\chi_1\gamma)}{dE_1^*}, \quad (\text{A6})$$

where $d\Gamma(P \rightarrow \bar{\chi}_0\chi_1\gamma)/dE_1^*$ denotes the corresponding differential partial decay width and where

$$\Gamma(P \rightarrow \gamma\gamma) = \frac{\alpha^2 m_\pi^3}{32\pi^3 F_P^2} \quad (\text{A7})$$

denotes the full partial width of P to photon pairs. This differential partial width is given by

$$\frac{d\Gamma(P \rightarrow \bar{\chi}_0\chi_1\gamma)}{dE_1^*} = \int \frac{d^3 p_0 d^3 p_\gamma d\Omega_1}{16(2\pi)^5 m_P} |\mathcal{M}(P \rightarrow \bar{\chi}_0\chi_1\gamma)|^2 \times \frac{p_1^*}{E_0 E_\gamma} \delta^4(p_\gamma + p_1 + p_0 - m_P),$$

where $|\mathcal{M}(P \rightarrow \bar{\chi}_0\chi_1\gamma)|^2$ is the squared matrix element for the decay process $P \rightarrow \bar{\chi}_0\chi_1\gamma$, summed over final-state-particle spins. For the case of an MDM operator structure, this squared matrix element is given by

$$|\mathcal{M}_{\text{MDM}}(P \rightarrow \bar{\chi}_0\chi_1\gamma)|^2 = \frac{8e^4}{\pi^4 F_P^2 \Lambda_m^2 q^4} [(2k_0 \cdot q)(k_1 \cdot q)(k_\gamma \cdot q)^2 + q^4(k_\gamma \cdot k_0)(k_\gamma \cdot k_1) - q^2(k_\gamma \cdot k_1)(k_\gamma \cdot q)(k_0 \cdot q) - m_0 m_1 q^2(k_\gamma \cdot q)^2 - q^2(k_\gamma \cdot k_0)(k_\gamma \cdot q)(k_1 \cdot q)], \quad (\text{A8})$$

where p , k_γ , k_0 , and k_1 respectively denote the four-momenta of P , γ , χ_0 , and χ_1 , and where q denotes the four-momentum of the virtual photon. By contrast, for the case of an EDM operator structure, this matrix element is given by

$$|\mathcal{M}_{\text{EDM}}(P \rightarrow \bar{\chi}_0\chi_1\gamma)|^2 = \frac{8e^4}{\pi^4 F_P^2 \Lambda_e^2 q^4} [(2k_0 \cdot q)(k_1 \cdot q)(k_\gamma \cdot q)^2 + q^4(k_\gamma \cdot k_0)(k_\gamma \cdot k_1) - q^2(k_\gamma \cdot k_1)(k_\gamma \cdot q)(k_0 \cdot q) + m_0 m_1 q^2(k_\gamma \cdot q)^2 - q^2(k_\gamma \cdot k_0)(k_\gamma \cdot q)(k_1 \cdot q)]. \quad (\text{A9})$$

We note that the factors of F_P in $\Gamma(P \rightarrow \gamma\gamma)$ and $d\Gamma(P \rightarrow \bar{\chi}_0\chi_1\gamma)/dE_1^*$ cancel in Eq. (A6), while $\text{BR}(P \rightarrow \gamma\gamma)$ is well approximated by its SM value.

APPENDIX B: THERMALLY AVERAGED CROSS-SECTIONS AND DECAY RATES

In this appendix, we provide expressions for the relevant thermally averaged quantities that appear in the collision terms Eq. (4.3).

In general, the thermally averaged decay width for a particle that decays via the two-body process $a \rightarrow b + c$,

where the species c and d in the final state are in thermal equilibrium with the SM radiation bath, is given by

$$\langle \Gamma_{a \rightarrow bc} \rangle = \frac{g_a}{n_a^{\text{eq}}} \int \frac{d^3 p_a}{(2\pi)^3} \frac{m_a}{E_a} e^{-E_a/T} \Gamma_{a \rightarrow bc}, \quad (\text{B1})$$

where g_a , n_a^{eq} , m_a , and $\Gamma_{a \rightarrow bc}$, respectively denote the number of internal degrees of freedom, and equilibrium number density, mass, and proper decay width of the initial-state particle a . Since the equilibrium number density is

$$n_a^{\text{eq}} = g_a \frac{m_a^2 T}{2\pi^2} K_2(m_a/T), \quad (\text{B2})$$

where $K_\nu(x)$ denotes the modified Bessel function of the second kind, one finds that

$$\langle \Gamma_{a \rightarrow bc} \rangle = \frac{K_1(m_a/T)}{K_2(m_a/T)} \Gamma_{a \rightarrow bc}. \quad (\text{B3})$$

The sole thermally averaged decay width that appears in Eq. (4.3) is the one for the decay process $\chi_1 \rightarrow \chi_0 \gamma$. The corresponding proper decay width from which this thermal average is derived is provided in Eq. (3.5).

The thermally averaged cross-section for the $2 \rightarrow 2$ scattering process of the form $a + b \rightarrow c + d$, where the species c and d in the final state are in thermal equilibrium with the SM radiation bath, can be written as [75]

$$\langle \sigma v \rangle = \frac{1}{n_a^{\text{eq}} n_b^{\text{eq}}} \int \frac{d^3 p_a}{(2\pi)^3} \frac{d^3 p_b}{(2\pi)^3} \sigma_{ab \rightarrow cd} v_{\text{rel}} e^{-(E_a + E_b)/T}, \quad (\text{B4})$$

where $\sigma_{ab \rightarrow cd}$ is the invariant cross-section for the scattering process and where

$$v_{\text{rel}} \equiv \frac{\sqrt{(p_a \cdot p_b)^2 - m_a^2 m_b^2}}{p_a \cdot p_b} \quad (\text{B5})$$

denotes the relative velocity of a and b in the background frame. To a very good approximation, when a and b are highly nonrelativistic, this relative velocity coincides with the Møller velocity and the expression in Eq. (B4) simplifies to [76]

$$\langle \sigma v \rangle = \frac{T}{32\pi^4 n_a^{\text{eq}} n_b^{\text{eq}}} \int_{s_{\text{min}}}^{\infty} ds (s - s_{\text{min}}) [s - (m_a - m_b)^2] \times \frac{\sigma_{ab \rightarrow cd}}{\sqrt{s}} K_1\left(\frac{\sqrt{s}}{T}\right), \quad (\text{B6})$$

where $s_{\text{min}} \equiv (m_a + m_b)^2$. For the case in which a and b are of the same species, this expression further reduces to

$$\langle \sigma v \rangle = \frac{T}{32\pi^4 (n_a^{\text{eq}})^2} \int_{4m_a^2}^{\infty} ds (s - 4m_a^2) \sqrt{s} \sigma_{ab \rightarrow cd} K_1\left(\frac{\sqrt{s}}{T}\right). \quad (\text{B7})$$

We now turn to discuss the cross-sections for the individual processes whose thermal averages appear in Eq. (4.3). For both the MDM and EDM operator structures, the cross-section for the s -channel co-annihilation process $\bar{\chi}_0 \chi_1 \rightarrow e^+ e^-$ can be written as

$$\sigma_{01 \rightarrow ee} = \frac{e^2}{6\pi \Lambda_{\mathcal{O}}^2} \frac{\sqrt{s - 4m_e^2} C^{(s)}}{s^{5/2} \sqrt{(s - s_{\text{min}}) [s - (m_0 - m_1)^2]}}, \quad (\text{B8})$$

where $s_{\text{min}} = (m_0 + m_1)^2$ and where the form of $C^{(s)}$ differs between the MDM and EDM cases. For the MDM case, this quantity takes the form

$$C^{(s)} = C_1^{(s)} + C_2^{(s)}, \quad (\text{B9})$$

where we have defined

$$\begin{aligned} C_1^{(s)} &\equiv 2m_e^2 [s^2 + (s_{\text{min}} + 4m_0 m_1) s - 2s_{\text{min}} (m_0 - m_1)^2], \\ C_2^{(s)} &\equiv s \left[s^2 + (s_{\text{min}} + 4m_0 m_1) s - \frac{7}{2} s_{\text{min}} (m_0 - m_1)^2 \right]. \end{aligned} \quad (\text{B10})$$

By contrast, for the EDM case, $C^{(s)}$ takes the form

$$C^{(s)} = (s + 2m_e^2) [s^2 + (s_{\text{min}} - 8m_0 m_1) s - 2s_{\text{min}} (m_0 - m_1)^2]. \quad (\text{B11})$$

For processes in which $\bar{\chi}_0$ and χ_1 co-annihilate into hadronic final states, we model the inclusive cross-section using the $R(s)$ function given in Ref. [77]. In particular, we take

$$\sigma_{01 \rightarrow \text{hadrons}} = \sigma_{01 \rightarrow \mu\mu} R(s) \Theta(s - 4m_\pi^2), \quad (\text{B12})$$

where $\Theta(x)$ denotes the Heaviside theta function and where m_π is the pion mass.

For both the MDM and EDM operator structures, the cross-section for the t -channel annihilation process $\bar{\chi}_0 \chi_0 \rightarrow \gamma\gamma$ can be written as

$$\sigma_{00 \rightarrow \gamma\gamma} = \frac{2}{\pi \Lambda_{\mathcal{O}}^4} \frac{C^{(t)}}{s(s - 4m_0^2)}, \quad (\text{B13})$$

where the form of $C^{(t)}$ differs between the MDM and EDM cases. In the MDM case, this quantity takes the form

$$C^{(t)} = \sum_{\ell=1}^5 C_\ell^{(t)} \quad (\text{B14})$$

where we have defined

$$\begin{aligned}
C_1^{(t)} &\equiv \frac{1}{3} [(m_0^2 - X_+)^3 - (m_0^2 - X_-)^3], \\
C_2^{(t)} &\equiv (m_0^2 - m_1^2)^4 \left(\frac{1}{m_1^2 - m_0^2 + X_+} - \frac{1}{m_1^2 - m_0^2 + X_-} \right), \\
C_3^{(t)} &\equiv \frac{(s + 2m_1^2 - 4m_0^2)}{2} [(m_0^2 - X_+)^2 - (m_0^2 - X_-)^2], \\
C_4^{(t)} &\equiv [6m_0^4 + 3m_1^4 - m_0^2(3s + 8m_1^2)](X_- - X_+), \\
C_5^{(t)} &\equiv \frac{1}{s + 2(m_1^2 - m_0^2)} \\
&\quad \times [8(m_0^2 - m_1^2)^4 - 4(m_0^2 - m_1^2)^2(2m_0^2 - m_1^2)s \\
&\quad + (m_0^4 - 4m_0^2m_1^2 - m_1^4)s^2] \log \left(\frac{X_+ - m_0^2 + m_1^2}{X_- - m_0^2 + m_1^2} \right), \tag{B15}
\end{aligned}$$

in terms of the quantities

$$X_{\pm} \equiv \frac{1}{2} \left(s \pm \sqrt{s(s - 4m_0^2)} \right). \tag{B16}$$

By contrast, for the EDM case, $C^{(t)}$ takes the form

$$C^{(t)} = \frac{1}{s - 2(m_0^2 - m_1^2)} \frac{X_+ - X_-}{m_1^2 s + (m_0^2 - m_1^2)^2} \sum_{\ell=1}^6 C_{\ell}^{(t)}, \tag{B17}$$

where we have defined

$$\begin{aligned}
C_1^{(t)} &\equiv 8(m_0^2 - m_1^2)^5, \\
C_2^{(t)} &\equiv -\frac{4}{3} (5m_0^2 - 6m_1^2)(m_0^2 - m_1^2)^3 s, \\
C_3^{(t)} &\equiv \frac{1}{3} (3m_0^6 - 19m_0^4 m_1^2 + 15m_0^2 m_1^4 + m_1^6) s^2, \\
C_4^{(t)} &\equiv \frac{1}{6} (m_0^4 + 4m_0^2 m_1^2 + 9m_1^4) s^3, \\
C_5^{(t)} &\equiv \frac{1}{6} m_1^2 s^4, \\
C_6^{(t)} &\equiv [8(m_0^2 - m_1^2)^4 - 4(m_0^2 - m_1^2)^2(2m_0^2 - m_1^2)s \\
&\quad + (m_0^4 - 4m_0^2 m_1^2 - m_1^4)s^2] \log \left(\frac{X_+ - m_0^2 + m_1^2}{X_- - m_0^2 + m_1^2} \right), \tag{B18}
\end{aligned}$$

and where X_{\pm} are once again given in Eq. (B16). The corresponding expressions for the t -channel annihilation process $\bar{\chi}_1 \chi_1 \rightarrow \gamma\gamma$ in the MDM and EDM cases are identical in form to the expression in Eq. (B13), but with $m_0 \leftrightarrow m_1$ exchanged in this cross-section formula and all ancillary expressions.

For both the MDM and EDM operator structures, the cross-section for the process $\chi_0 e^{\pm} \rightarrow \chi_1 e^{\pm}$, in which a χ_0 particle up-scatters into a χ_1 particle off an electron or positron, can be written as

$$\sigma_{0e \rightarrow 1e} = \frac{e^2 C^{(\text{sc})}}{2\pi \Lambda_{\mathcal{O}}^2 s (s - s_{\min}) [s - (m_0 - m_e)^2]} \tag{B19}$$

where $s_{\min} = (m_0 + m_e)^2$ and where the form of $C^{(\text{sc})}$ differs between the MDM and EDM cases. In the MDM case, this quantity takes the form

$$\begin{aligned}
C^{(\text{sc})} &= C_1^{(\text{sc})} [(m_0 + m_1)^2 + 2m_e^2 - 4s] \\
&\quad + s[2s^2 - 4m_e^2(s + m_0 m_1) + 2m_e^4 + m_0^4 + m_1^4 \\
&\quad - C_3^{(\text{sc})}] \log \left(\frac{C_3^{(\text{sc})} - C_1^{(\text{sc})} - C_2^{(\text{sc})}}{C_1^{(\text{sc})} - C_2^{(\text{sc})} + C_3^{(\text{sc})}} \right), \tag{B20}
\end{aligned}$$

where we have defined

$$\begin{aligned}
C_1^{(\text{sc})} &\equiv \sqrt{(s - s_{\min}) [s - (m_e - m_0)^2]} \\
&\quad \times \sqrt{[s - (m_e + m_1)^2] [s - (m_e - m_1)^2]}, \\
C_2^{(\text{sc})} &\equiv (s + m_0^2 - m_e^2)(s + m_1^2 - m_e^2), \\
C_3^{(\text{sc})} &\equiv 2s(m_0^2 + m_1^2). \tag{B21}
\end{aligned}$$

By contrast, in the EDM case, $C^{(\text{sc})}$ takes the form

$$\begin{aligned}
C^{(\text{sc})} &= C_1^{(\text{sc})} [(m_0 - m_1)^2 + 2m_e^2 - 4s] \\
&\quad + s[2s^2 - 4m_e^2(s - m_0 m_1) + 2m_e^4 + m_0^4 + m_1^4 \\
&\quad - C_3^{(\text{sc})}] \log \left(\frac{C_3^{(\text{sc})} - C_1^{(\text{sc})} - C_2^{(\text{sc})}}{C_1^{(\text{sc})} - C_2^{(\text{sc})} + C_3^{(\text{sc})}} \right). \tag{B22}
\end{aligned}$$

The cross-section for the corresponding down-scattering process $\chi_1 e^{\pm} \rightarrow \chi_0 e^{\pm}$ may be obtained simply by exchanging $m_0 \leftrightarrow m_1$ everywhere in Eq. (B19) and all ancillary expressions. Moreover, the corresponding expression for the up-scattering of χ_0 into χ_1 off any other electrically charged SM particle may be obtained by replacing m_e with the mass of that SM particle everywhere in Eq. (B19) and all ancillary expressions.

APPENDIX C: PROSPECTS FOR DETECTION AT SHiP

SHiP [27,28] is a proposed experiment at CERN in which a 400-GeV proton beam will collide with a fixed target consisting of tungsten and a titanium-zirconium-doped molybdenum alloy. SHiP probes lifetimes comparable to FASER and FASER2 and so it is useful to compare the reach of these experiments.

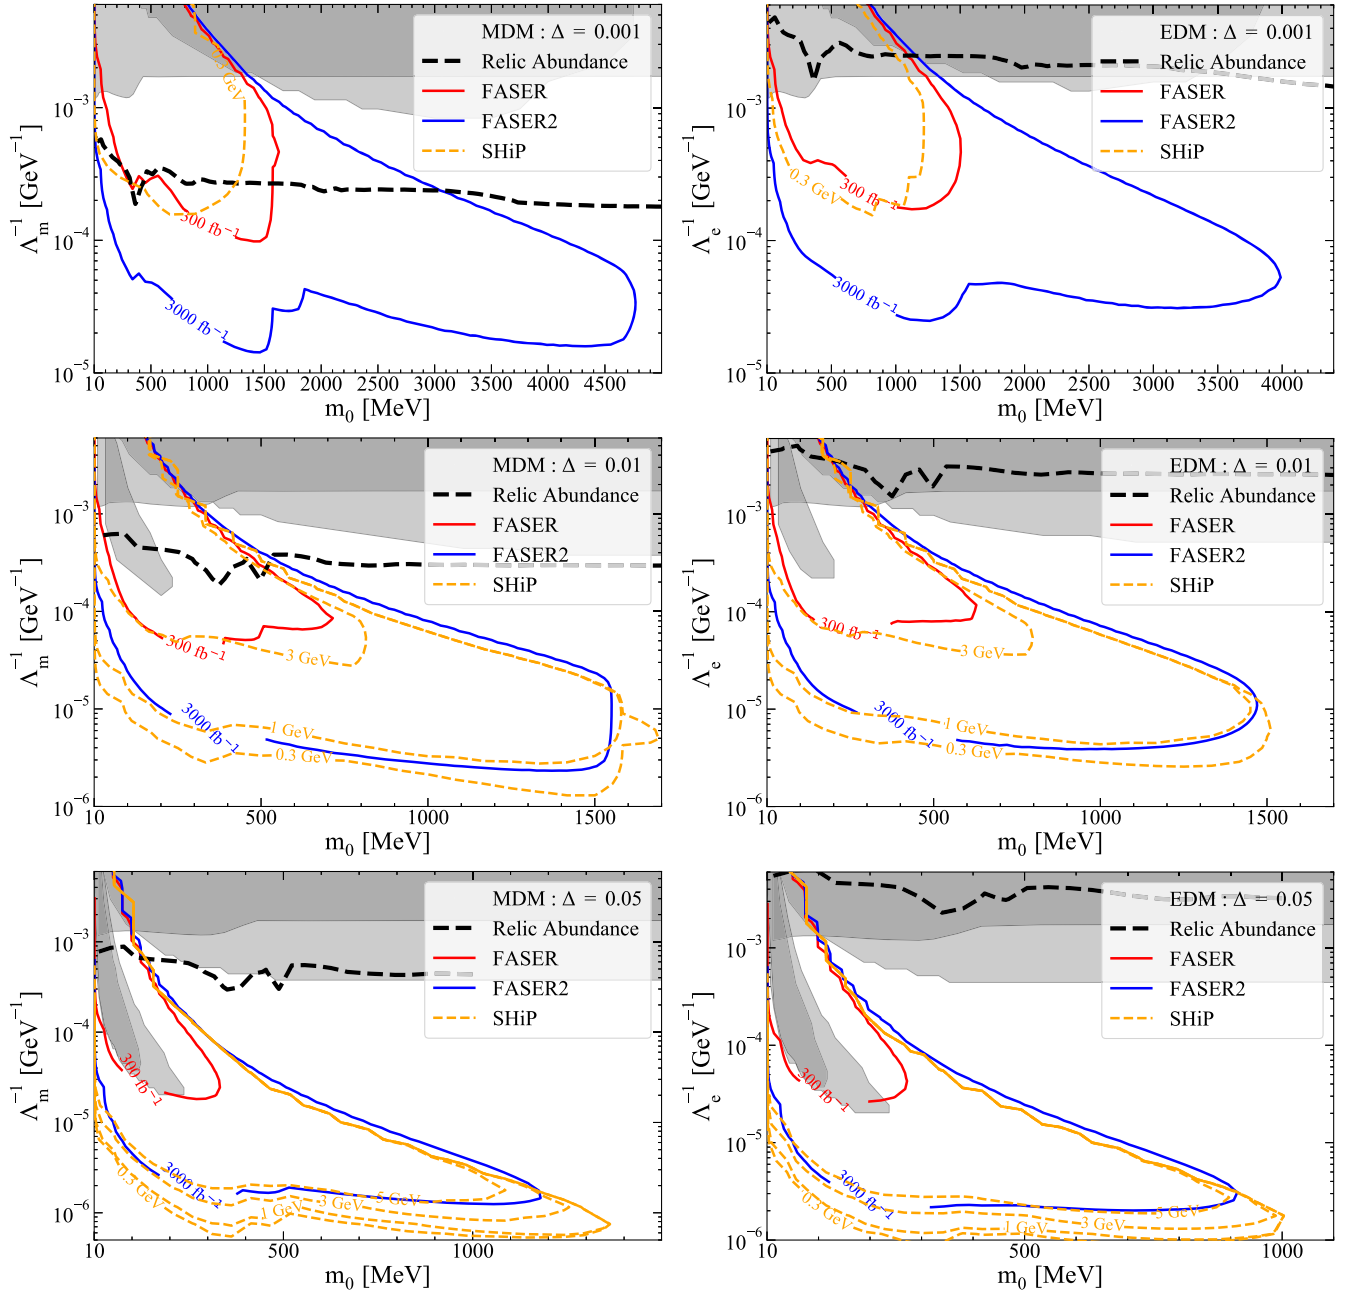


FIG. 5. Similar to Fig. 3, but including contours representing the discovery reach of SHiP (dashed orange curves) for different monophoton energy thresholds. Each such contour corresponds to the observation of three signal events for a total of 2×10^{20} protons on target. Contours are shown for the thresholds $E_\gamma \geq 0.3$ GeV, $E_\gamma \geq 1$ GeV, $E_\gamma \geq 3$ GeV, and $E_\gamma \geq 5$ GeV. We observe that SHiP has a comparable reach to FASER2 for large Δ , even for larger detection thresholds, but has far less sensitivity than FASER2 for small Δ , even for a detection threshold $E_\gamma \geq 0.3$ GeV.

As described in Ref. [78], SHiP is proposed to consist of a decay volume shaped as a 50-meter-long pyramidal frustum with an upstream face 45 meters from the target with dimensions $5 \text{ m} \times 2 \text{ m}$ and a downstream face with dimensions $11 \text{ m} \times 6 \text{ m}$. The SHiP detector is expected to be able to detect photons with energies $E_\gamma \gtrsim 300 \text{ MeV}$ [79]. However, resolving the monophoton signal that results from $\chi_1 \rightarrow \chi_0 \gamma$ decay from background may require a

somewhat higher photon-energy threshold. For example, in assessing the prospects for detecting a monophoton signature of dark-photon decay in the context of a dark-axion-portal model, the authors of Ref. [80] adopted a threshold $E_\gamma > 2 \text{ GeV}$ based on comparisons with the process in which the decay of the dark photon yields a e^+e^- pair rather than a photon. However, until a detailed monophoton study for SHiP is performed, it remains

unclear precisely what threshold should be adopted in searches of this sort. Thus, in what follows, we consider a range of possible choices for this threshold.

The procedure we use for determining the expected number of signal events at SHiP for any given combination of m_0 , Δ , and $\Lambda_{\mathcal{O}}$ is analogous to the procedure described in Sec. VB that we use in determining the expected number of events at other fixed-target experiments. We calculate the differential cross-section for π^0 production due to the scattering of protons in the beam with a molybdenum target using the BDNMC code package. We then scale this cross-section by $\mathcal{N}_M^{(\text{POT})}$ as in Eq. (5.3) to obtain differential production cross-sections for heavier neutral mesons. The $\mathcal{N}_M^{(\text{POT})}$ for all relevant meson species except the vector meson $\psi(2S)$ were provided in Ref. [31]. We calculate the number of $\psi(2S)$ mesons produced per proton using PYTHIA 8.2 and find that $\mathcal{N}_{\psi(2S)}^{(\text{POT})} = 10^{-6}$. We then determine the expected number of χ_1 decays per proton on target within the decay volume of the SHiP detector from the branching fractions given in Appendix A, assuming the detector geometry described above.

In Fig. 5, we show contours (orange dashed curves) of the discovery reach of the SHiP experiment within our parameter-space region of interest for 2×10^{20} total protons

on target—the number expected over roughly 10 years of operation [28]. Since this discovery reach depends sensitively on the threshold energy for a single photon, we plot contours for several different values of this threshold: $E_\gamma \geq 0.3$ GeV, 1 GeV, 3 GeV, and 5 GeV. As in Fig. 3, the panels shown in the left and right columns correspond to the choice of MDM and EDM operators, respectively, while the panels shown in different rows correspond to different values of Δ . Contours indicating the reach of FASER with an integrated luminosity of 300 fb^{-1} (red curves) and the reach of FASER2 with 3000 fb^{-1} (blue curves) are also shown for comparison.

We observe from Fig. 5 that for small values of Δ , the SHiP detector has little sensitivity within our parameter-space region of interest. This is primarily a reflection of the fact that for small mass splittings, the lab-frame energy of the photon produced by χ_1 decay typically lies below the threshold for detection—even if that threshold is as low as 300 MeV. By contrast, for large values of Δ , a larger fraction of the photons have energies that exceed the detection threshold and the discovery reach of SHiP is far greater. For example, for $\Delta = 0.05$, the reach of SHiP and FASER2 is very roughly similar for comparable running times.

-
- [1] A. Boveia and C. Doglioni, Dark matter searches at colliders, *Annu. Rev. Nucl. Part. Sci.* **68**, 429 (2018).
 - [2] L. Lee, C. Ohm, A. Soffer, and T.-T. Yu, Collider searches for long-lived particles beyond the standard model, *Prog. Part. Nucl. Phys.* **106**, 210 (2019); **122**, 103912(E) (2022).
 - [3] J. Alimena *et al.*, Searching for long-lived particles beyond the Standard Model at the Large Hadron Collider, *J. Phys. G* **47**, 090501 (2020).
 - [4] J. L. Feng, I. Galon, F. Kling, and S. Trojanowski, Forward Search Experiment at the LHC, *Phys. Rev. D* **97**, 035001 (2018).
 - [5] A. Ariga *et al.* (FASER Collaboration), FASER’s physics reach for long-lived particles, *Phys. Rev. D* **99**, 095011 (2019).
 - [6] H. Abreu *et al.* (FASER Collaboration), The FASER detector, [arXiv:2207.11427](https://arxiv.org/abs/2207.11427).
 - [7] R. Mammen Abraham *et al.*, Forward Physics Facility—Snowmass 2021 Letter of Interest, <https://zenodo.org/record/4059893>.
 - [8] L. A. Anchordoqui *et al.*, The Forward Physics Facility: Sites, experiments, and physics potential, *Phys. Rep.* **968**, 1 (2022).
 - [9] J. L. Feng *et al.*, The forward physics facility at the high-luminosity LHC, *J. Phys. G* **50**, 030501 (2023).
 - [10] T. Han and R. Hempfling, Messenger sneutrinos as cold dark matter, *Phys. Lett. B* **415**, 161 (1997).
 - [11] L. J. Hall, T. Moroi, and H. Murayama, Sneutrino cold dark matter with lepton number violation, *Phys. Lett. B* **424**, 305 (1998).
 - [12] D. Tucker-Smith and N. Weiner, Inelastic dark matter, *Phys. Rev. D* **64**, 043502 (2001).
 - [13] D. Barducci, E. Bertuzzo, M. Taoso, and C. Toni, Probing right-handed neutrinos dipole operators, *J. High Energy Phys.* **03** (2023) 239.
 - [14] F. Kling and P. Quilés, ALP searches at the LHC: FASER as a light-shining-through-walls experiment, *Phys. Rev. D* **106**, 055036 (2022).
 - [15] H. K. Dreiner, D. Köhler, S. Nangia, and Z. S. Wang, Searching for a single photon from lightest neutralino decays in R-parity-violating supersymmetry at FASER, *J. High Energy Phys.* **02** (2023) 120.
 - [16] K. R. Dienes and B. Thomas, More is different: Non-minimal dark sectors and their implications for particle physics, astrophysics, and cosmology—13 take-away lessons for snowmass 2021, in 2022 Snowmass Summer Study, [arXiv:2203.17258](https://arxiv.org/abs/2203.17258).
 - [17] K. R. Dienes, J. Kumar, B. Thomas, and D. Yaylali, Dark-Matter Decay as a Complementary Probe of Multi-component Dark Sectors, *Phys. Rev. Lett.* **114**, 051301 (2015).
 - [18] K. R. Dienes, J. Kumar, B. Thomas, and D. Yaylali, Off-diagonal dark-matter phenomenology: Exploring enhanced

- complementarity relations in nonminimal dark sectors, *Phys. Rev. D* **96**, 115009 (2017).
- [19] K. R. Dienes, D. Kim, H. Song, S. Su, B. Thomas, and D. Yaylali, Nonminimal dark sectors: Mediator-induced decay chains and multijet collider signatures, *Phys. Rev. D* **101**, 075024 (2020).
- [20] K. R. Dienes, D. Kim, T. T. Leininger, and B. Thomas, Sequential displaced vertices (“Tumblers”): A novel collider signature for long-lived particles, *Phys. Rev. D* **106**, 095012 (2022).
- [21] K. R. Dienes and B. Thomas, Dynamical dark matter: I. Theoretical overview, *Phys. Rev. D* **85**, 083523 (2012).
- [22] K. R. Dienes and B. Thomas, Dynamical dark matter: II. An explicit model, *Phys. Rev. D* **85**, 083524 (2012).
- [23] K. R. Dienes and B. Thomas, Phenomenological constraints on axion models of dynamical dark matter, *Phys. Rev. D* **86**, 055013 (2012).
- [24] S. Chang, N. Weiner, and I. Yavin, Magnetic inelastic dark matter, *Phys. Rev. D* **82**, 125011 (2010).
- [25] E. Izaguirre, G. Krnjaic, and B. Shuve, Discovering inelastic thermal-relic dark matter at colliders, *Phys. Rev. D* **93**, 063523 (2016).
- [26] E. Masso, S. Mohanty, and S. Rao, Dipolar dark matter, *Phys. Rev. D* **80**, 036009 (2009).
- [27] W. Bonivento *et al.*, Proposal to search for heavy neutral leptons at the SPS, [arXiv:1310.1762](https://arxiv.org/abs/1310.1762).
- [28] M. Anelli *et al.* (SHiP Collaboration), A facility to search for hidden particles (SHiP) at the CERN SPS, [arXiv:1504.04956](https://arxiv.org/abs/1504.04956).
- [29] X. Chu, J. Pradler, and L. Semmelrock, Light dark states with electromagnetic form factors, *Phys. Rev. D* **99**, 015040 (2019).
- [30] F. Kling, J.-L. Kuo, S. Trojanowski, and Y.-D. Tsai, FLArE up dark sectors with EM form factors at the LHC forward physics facility, *Nucl. Phys.* **B987**, 116103 (2023).
- [31] X. Chu, J.-L. Kuo, and J. Pradler, Dark sector-photon interactions in proton-beam experiments, *Phys. Rev. D* **101**, 075035 (2020).
- [32] N. Aghanim *et al.* (Planck Collaboration), Planck 2018 results. VI. Cosmological parameters, *Astron. Astrophys.* **641**, A6 (2020); **652**, C4(E) (2021).
- [33] K. Jodłowski, F. Kling, L. Roszkowski, and S. Trojanowski, Extending the reach of FASER, MATHUSLA, and SHiP towards smaller lifetimes using secondary particle production, *Phys. Rev. D* **101**, 095020 (2020).
- [34] F. Kling and S. Trojanowski, Forward experiment sensitivity estimator for the LHC and future hadron colliders, *Phys. Rev. D* **104**, 035012 (2021).
- [35] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations, *J. High Energy Phys.* **07** (2014) 079.
- [36] K. Jodłowski and S. Trojanowski, Neutrino beam-dump experiment with FASER at the LHC, *J. High Energy Phys.* **05** (2021) 191.
- [37] J. Arakawa, J.L. Feng, A. Ismail, F. Kling, and M. Waterbury, Neutrino detection without neutrino detectors: Discovering collider neutrinos at FASER with electronic signals only, *Phys. Rev. D* **106**, 052011 (2022).
- [38] A. Ariga *et al.* (FASER Collaboration), Technical proposal for FASER: Forward search experiment at the LHC, [arXiv:1812.09139](https://arxiv.org/abs/1812.09139).
- [39] 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).
- [40] T. Abe *et al.* (Belle-II Collaboration), Belle II technical design report, [arXiv:1011.0352](https://arxiv.org/abs/1011.0352).
- [41] J. D. Bjorken, S. Ecklund, W. R. Nelson, A. Abashian, C. Church, B. Lu, L. W. Mo, T. A. Nunamaker, and P. Rassmann, Search for neutral metastable penetrating particles produced in the SLAC beam dump, *Phys. Rev. D* **38**, 3375 (1988).
- [42] E. M. Riordan *et al.*, Search for Short Lived Axions in an Electron Beam Dump Experiment, *Phys. Rev. Lett.* **59**, 755 (1987).
- [43] M. Davier and H. Nguyen Ngoc, An unambiguous search for a light Higgs boson, *Phys. Lett. B* **229**, 150 (1989).
- [44] A. Konaka *et al.*, Search for Neutral Particles in Electron Beam Dump Experiment, *Phys. Rev. Lett.* **57**, 659 (1986).
- [45] S. Andreas, C. Niebuhr, and A. Ringwald, New limits on hidden photons from past electron beam dumps, *Phys. Rev. D* **86**, 095019 (2012).
- [46] B. Batell, R. Essig, and Z. Surujon, Strong Constraints on Sub-GeV Dark Sectors from SLAC Beam Dump E137, *Phys. Rev. Lett.* **113**, 171802 (2014).
- [47] J. D. Bjorken, R. Essig, P. Schuster, and N. Toro, New fixed-target experiments to search for dark gauge forces, *Phys. Rev. D* **80**, 075018 (2009).
- [48] K. J. Kim and Y.-S. Tsai, Improved Weizsacker-Williams method and its application to lepton and W boson pair production, *Phys. Rev. D* **8**, 3109 (1973).
- [49] J. Blumlein and J. Brunner, New exclusion limits for dark gauge forces from beam-dump data, *Phys. Lett. B* **701**, 155 (2011).
- [50] 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).
- [51] P. Vilain *et al.* (CHARM-II Collaboration), Precision measurement of electroweak parameters from the scattering of muon-neutrinos on electrons, *Phys. Lett. B* **335**, 246 (1994).
- [52] S. N. Gninenko, Constraints on sub-GeV hidden sector gauge bosons from a search for heavy neutrino decays, *Phys. Lett. B* **713**, 244 (2012).
- [53] V. V. Ammosov *et al.* (French-Soviet Collaboration), $\pi^+\pi^-$ meson and proton production in inclusive and seminclusive processes in pp interactions at 69-GeV/c. 2., *Nuovo Cimento A* **40**, 237 (1977).
- [54] T. Sjöstrand, S. Mrenna, and P. Z. Skands, PYTHIA 6.4 physics and manual, *J. High Energy Phys.* **05** (2006) 026.
- [55] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C. O. Rasmussen, and P. Z. Skands, An introduction to PYTHIA 8.2, *Comput. Phys. Commun.* **191**, 159 (2015).
- [56] 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).
- [57] K. De Winter *et al.* (CHARM-II Collaboration), A detector for the study of neutrino-electron scattering, *Nucl. Instrum. Methods Phys. Res., Sect. A* **278**, 670 (1989).
- [58] M. Bonesini, A. Marchionni, F. Pietropaolo, and T. Tabarelli de Fatis, On particle production for high-energy neutrino beams, *Eur. Phys. J. C* **20**, 13 (2001).
- [59] 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).
- [60] F. Bergsma *et al.* (CHARM Collaboration), A search for decays of heavy neutrinos, *Phys. Lett. B* **128**, 361 (1983).
- [61] S. Mohanty and S. Rao, Detecting dipolar dark matter in beam dump experiments, [arXiv:1506.06462](https://arxiv.org/abs/1506.06462).
- [62] A. A. Aguilar-Arevalo *et al.* (MiniBooNE Collaboration), The MiniBooNE detector, *Nucl. Instrum. Methods Phys. Res., Sect. A* **599**, 28 (2009).
- [63] C. Athanassopoulos *et al.* (LSND Collaboration), The liquid scintillator neutrino detector and LAMPF neutrino source, *Nucl. Instrum. Methods Phys. Res., Sect. A* **388**, 149 (1997).
- [64] J.-F. Fortin and T. M. P. Tait, Collider constraints on dipole-interacting dark matter, *Phys. Rev. D* **85**, 063506 (2012).
- [65] R. Bartels, D. Gaggero, and C. Weniger, Prospects for indirect dark matter searches with MeV photons, *J. Cosmol. Astropart. Phys.* **05** (2017) 001.
- [66] K. K. Boddy and J. Kumar, Indirect detection of dark matter using MeV-range gamma-ray telescopes, *Phys. Rev. D* **92**, 023533 (2015).
- [67] P. A. R. Ade *et al.* (Planck Collaboration), Planck 2015 results. XIII. Cosmological parameters, *Astron. Astrophys.* **594**, A13 (2016).
- [68] N. Sabti, J. Alvey, M. Escudero, M. Fairbairn, and D. Blas, Refined bounds on MeV-scale thermal dark sectors from BBN and the CMB, *J. Cosmol. Astropart. Phys.* **01** (2020) 004.
- [69] 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).
- [70] W. Baldini *et al.*, SHADOWS (search for hidden and dark objects with the SPS), [arXiv:2110.08025](https://arxiv.org/abs/2110.08025).
- [71] M. Bauer, O. Brandt, L. Lee, and C. Ohm, ANUBIS: Proposal to search for long-lived neutral particles in CERN service shafts, [arXiv:1909.13022](https://arxiv.org/abs/1909.13022).
- [72] G. Aielli *et al.*, Expression of interest for the CODEX-b detector, *Eur. Phys. J. C* **80**, 1177 (2020).
- [73] B. Acharya *et al.*, MoEDAL-MAPP, an LHC dedicated detector search facility, in 2022 Snowmass Summer Study, [arXiv:2209.03988](https://arxiv.org/abs/2209.03988).
- [74] H. Lubatti *et al.*, MATHUSLA: A detector proposal to explore the lifetime frontier at the HL-LHC, [arXiv:1901.04040](https://arxiv.org/abs/1901.04040).
- [75] M. Cannoni, Relativistic $\langle\sigma v_{\text{rel}}\rangle$ in the calculation of relics abundances: A closer look, *Phys. Rev. D* **89**, 103533 (2014).
- [76] P. Gondolo and G. Gelmini, Cosmic abundances of stable particles: Improved analysis, *Nucl. Phys.* **B360**, 145 (1991).
- [77] R. L. Workman *et al.* (Particle Data Group), Review of particle physics, *Prog. Theor. Exp. Phys.* **2022**, 083C01 (2022).
- [78] SHiP Collaboration, SHiP experiment—comprehensive design study report, Technical Reports No. CERN-SPSC-2019-049, No. SPSC-SR-263, CERN, Geneva, 2019, <https://cds.cern.ch/record/2704147>.
- [79] A. Golutvin (private communication).
- [80] P. deNiverville and H.-S. Lee, Implications of the dark axion portal for SHiP and FASER and the advantages of mono-photon signals, *Phys. Rev. D* **100**, 055017 (2019).