

Dark matter freeze-in and small-scale observables: Novel mass bounds and viable particle candidates

Francesco D'Eramo^{1,2,*}, Alessandro Lenoci^{3,4,†} and Ariane Dekker^{5,‡}

¹*Dipartimento di Fisica e Astronomia, Università degli Studi di Padova,
Via Marzolo 8, 35131 Padova, Italy*

²*Istituto Nazionale di Fisica Nucleare (INFN), Sezione di Padova, Via Marzolo 8, 35131 Padova, Italy*

³*Racah Institute of Physics, The Hebrew University, 91904, Jerusalem, Israel*

⁴*Laboratory for Elementary Particle Physics, Cornell University, Ithaca, New York 14853, USA*

⁵*Kavli Institute for Cosmological Physics, University of Chicago, Chicago, Illinois 60637, USA*



(Received 24 June 2025; accepted 11 November 2025; published 8 December 2025)

The suppression of cosmological structure at small scales is a key signature of dark matter (DM) produced via freeze-in in the low-mass regime. We present a comprehensive analysis of its impact, incorporating recent constraints from Milky Way satellite counts, strong gravitational lensing with James Webb Space Telescope (JWST) data, and the Lyman- α forest. We adopt a general strategy to translate existing warm dark matter (WDM) bounds into lower mass limits for a broad class of DM candidates characterized by quasithermal phase space distributions. The benefits of this approach include computational efficiency and the ability to explore a wide range of models. We derive model-independent bounds for DM produced via two-body decays, scatterings, and three-body decays, and apply the framework to concrete scenarios such as the Higgs portal, sterile neutrinos, axionlike particles, and the dark photon portal. Results from specific models confirm the validity of the model-independent analysis.

DOI: [10.1103/j62q-cvkr](https://doi.org/10.1103/j62q-cvkr)

I. INTRODUCTION

Approximately half a century after its theoretical formulation, the Standard Model (SM) of particle physics continues to be confirmed by experiments at accelerators and colliders. Nevertheless, it leaves several fundamental questions about the Universe unanswered. Chief among these is the nature of dark matter (DM), whose particle identity remains unknown [1–5]. None of the numerous proposed theoretical candidates has yet received experimental confirmation.

In this paper, we focus on the *freeze-in paradigm*, in which the hypothetical DM particle interacts so feebly with SM degrees of freedom that it never attains thermal equilibrium throughout cosmic history. Despite their feebleness, these interactions can still generate DM particles through occasional decays and scatterings involving constituents of the primordial thermal bath. Once produced, the

DM particles redshift with the expansion of the Universe and persist to the present day.

As is well known, when DM interactions with the bath are renormalizable, freeze-in production is most efficient when the temperature is comparable to the mass of the heaviest particle involved in the process [6–9]. This implies that the DM relic abundance can be predicted solely in terms of particle physics parameters such as masses and couplings that could, in principle, be tested in laboratory experiments, without sensitivity to the unknown details of inflationary reheating. This feature is described by saying that freeze-in production is “IR-dominated.” This is undoubtedly a strength of the framework. It is also important to recognize how the feeble interactions render it extremely challenging to test experimentally. Promising avenues include displaced events at colliders [10–19], direct detection [20–24], and cosmological signatures imprinted on small-scale structure [25–40].

The main subject of this paper is the latter effect. Freeze-in models can feature large free-streaming lengths, which erase perturbations below a characteristic scale and consequently suppress the formation of small-scale structure. It is therefore essential to employ probes that are sensitive to the DM distribution at these scales. Several observational strategies have been used to constrain DM properties. In particular, observations of ultrafaint dwarf galaxies, the Lyman- α forest, strong gravitational lensing, and stellar

* Contact author: francesco.deramo@pd.infn.it

† Contact author: alessandro.lenoci@mail.huji.ac.il

‡ Contact author: ahdekker@uchicago.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.

streams have all been successfully employed to place limits on models that suppress or enhance structure on small scales (see, e.g., the reviews provided by Refs. [41,42] and the references cited in Sec. II). It is important to recognize that these constraints are complementary, as they rely on different observational techniques and are affected by distinct systematic uncertainties. Our work examines the implications of this collection of small-scale probes.

The suppression of small-scale structure is a known consequence of warm dark matter (WDM), whose cosmological effects have been extensively studied. This similarity with freeze-in DM allows us to construct a mapping between the two frameworks that provides multiple advantages. Most notably, it offers computational efficiency by avoiding the numerically intensive calculation of the matter power spectrum for each freeze-in model. In addition, it enables us to reinterpret existing mass bounds from the minimal WDM scenario. To implement this strategy, we solve the momentum-space Boltzmann equation to determine the DM phase space distribution (PSD) resulting from freeze-in production. We then constrain freeze-in DM by imposing limits on its *warmness*, quantified by the second moment of the PSD.

We begin with a model-independent analysis of DM production via three different channels: two-body decays, binary scatterings, and three-body decays. Throughout, we assume that masses and couplings are chosen such that the DM relic abundance matches the observed value; to this end, we solve the relevant Boltzmann equation and compute the relic density self-consistently. Our general treatment assumes constant transition amplitudes, which is strictly valid only for two-body decays. In the case of scatterings and three-body decays, the matrix elements typically exhibit a nontrivial dependence on the momenta of the external particles. However, since freeze-in production is IR dominated and the relevant dynamics occur at a characteristic energy scale, this approximation remains well motivated. We confirm this expectation by analyzing several explicit microscopic realizations, where we compute the exact transition amplitudes and use them as input to the momentum-space Boltzmann equation.

The paper is structured as follows. In Sec. II, we review the most up-to-date small-scale structure bounds on WDM and discuss the associated uncertainties and caveats. Section III introduces the concept of a *quasithermal* PSD and outlines a general strategy to translate WDM limits into bounds applicable to this broader class of candidates. In Sec. IV, we focus on freeze-in production and derive model-independent lower bounds on the DM mass. Section V presents explicit realizations of freeze-in via two-body decays in Higgs portal scenarios and in models involving sterile neutrinos. Section VI explores axionlike particles (ALPs) coupled to Standard Model fermions, while Sec. VII analyzes the dark photon portal. For each model, we compute the resulting PSD, extract the

corresponding mass bounds, and verify consistency with the model-independent analysis of Sec. IV. Technical details are deferred to the Appendixes, and our conclusions are summarized in Sec. VIII. The main results of this work—both the model-independent analysis and the model-specific mass bounds—are compactly and clearly illustrated in the summary plot of Fig. 7.

II. WARM DARK MATTER MASS BOUNDS

Small-scale tests of WDM are subject to significant observational and modeling uncertainties, making independent constraints from multiple probes essential. In this section, we review current WDM mass bounds, present an updated limit from Milky Way satellite counts with explicit assumptions (see Fig. 1), and briefly discuss the prospects of combining probes in joint analyses. A summary of the mass bounds is shown in Fig. 2.

Updated Milky Way satellite counts. The number of satellite galaxies in the MW is expected to be reduced in the case of WDM, as the formation of low-mass subhalos capable of hosting galaxies is suppressed relative to CDM. Therefore, by comparing the predicted number of satellite galaxies in a given WDM model to the observed count (N_{obs}), one can place strong constraints on WDM [43–49]. Uncertainties in these constraints arise, for instance, from the number of satellites associated with the large Magellanic cloud, the halo-galaxy connection [41], and the mass of the MW halo, which is estimated to lie in the range $(0.5 - 1.5) \times 10^{12} M_{\odot}$ [50–56].

Conservative limits can be obtained by assuming that all subhalos are capable of forming galaxies and by considering a wide range of MW halo masses. By combining data from the Dark Energy Survey (DES) and the Sloan Digital Sky Survey (SDSS) with $N_{\text{obs}} = 124^{+40}_{-27}$, Ref. [45] finds




FIG. 1. The constraints on m_{WDM} at 95% CL, as a function of the number of observable satellite galaxies derived following the framework of Ref. [43]. The dotted lines show the limits adopted in this work.




FIG. 2. WDM mass bounds from small-scale observations: the blue identifies the stringent bound from MW satellites; the orange from strong lensing due to JWST observations; the green from Lyman- α forests; the red the conservative bound from MW satellites; the violet from stellar streams; the brown from UV luminosity. See Sec. II for details.

$m_{\text{WDM}} > 2.02$ keV (95% CL) after marginalizing over the MW halo mass. Moreover, using updated data from DES and Pan-STARRS1 with $N_{\text{obs}} = 270$, Ref. [43] obtains conservative bounds of $m_{\text{WDM}} > 3.6$ – 5.1 keV (95% CL), which vary only weakly with the MW halo mass. After correcting for anisotropic satellite distributions, N_{obs} was updated to 220 ± 50 [57], which has only a minor impact on the conservative bounds. Incorporating a galaxy formation model leads to significantly stronger constraints, as it further reduces the total number of satellites. For instance, Ref. [44] derives a bound of $m_{\text{WDM}} > 6.5$ keV (95% CL) using an empirical galaxy–halo model and updated DES and Pan-STARRS1 data with $N_{\text{obs}} = 220 \pm 50$.

When setting bounds using a galaxy–halo model, the uncertainty in the observed total number of satellites becomes important. In the CDM scenario, up to ~ 150 satellites are predicted to remain undetected, and approximately one-quarter of these are expected to be associated with the large Magellanic cloud (LMC) [57]. Different studies adopt varying assumptions for N_{obs} , leading to different constraints on the WDM mass. Following the framework of Ref. [43], we revisit WDM bounds for different values of N_{obs} . We adopt a Milky Way halo mass of $M_{200}^{\text{DM}} = 0.88 \times 10^{12} M_{\odot}$, based on the weighted average of 20 individual estimates [58]. Figure 1 shows the derived 95% CL limits on the WDM mass under two scenarios: a conservative case assuming all subhalos host galaxies (violet, $M_{\text{thresh}} = \infty$), and a model with a galaxy formation threshold (green, $M_{\text{thresh}} = 10^8 M_{\odot}$). As expected, the limits are more sensitive to N_{obs} when a galaxy formation threshold is applied, while in the conservative case, the dependence on N_{obs} is weak.

Throughout this work, we adopt the galaxy formation threshold scenario with $N_{\text{obs}} = 220$, which yields a lower bound of $m_{\text{WDM}} > 6.8$ keV (95% CL). For comparison, we also consider a conservative bound of $m_{\text{WDM}} > 3.9$ keV (95% CL), assuming all subhalos host galaxies and

adopting $N_{\text{obs}} = 183$, which excludes the estimated number of satellites associated with the LMC.

Strong gravitational lensing. DM subhalos within lens galaxies and halos along the line of sight can perturb strong gravitational lens images. These perturbations serve as probes of low-mass halos and allow constraints on WDM models based purely on the gravitational potential, thereby avoiding uncertainties related to baryonic physics. Systematic uncertainties arise from the modeling of the lens, source, and subhalo structures, as well as from flux ratio measurements. These are expected to improve with upcoming surveys and high-resolution follow-up observations [59]. Surface brightness distortions from lensed quasar host galaxies, observed as arcs, have proven to be effective probes [60]. Using seven observations of multiply imaged quasars, Ref. [61] excludes $m_{\text{WDM}} < 5.58$ keV at 95% CL. Another method relies on the relative magnifications of lensed images [62]. Based on Hubble Space Telescope observations of eight quadruply imaged quasars, Ref. [44] reports $m_{\text{WDM}} > 4.9$ keV at 95% confidence. More recently, James Webb Space Telescope (JWST) observations of 31 lensed quasars strengthen this constraint to $m_{\text{WDM}} > 6.1$ keV at posterior odds of 10:1 [63].

Lyman- α forest. The Lyman- α forest consists of absorption features in the spectra of high-redshift quasars, caused by neutral hydrogen in the intergalactic medium (IGM) along the line of sight. These absorption lines trace the underlying matter density field and can therefore be used to constrain WDM by comparing high-resolution spectral data with predictions from hydrodynamical simulations [64–66]. These simulations carry significant uncertainties, particularly related to the thermal history of the IGM [67]. Using 1080 high-resolution cosmological hydrodynamical simulations along with quasar spectra from the Keck Observatory and the Very Large Telescope, Ref. [68] derives a lower limit of $m_{\text{WDM}} > 3.1$ keV (95% CL). More recently, Ref. [69] sets a stronger constraint of $m_{\text{WDM}} > 5.7$ keV (95% CL) using high-resolution spectra of distant quasars obtained with the HIRES and UVES spectrographs. This improvement is attributed to a larger dataset and reduced observational uncertainties, enabling the probing of smaller cosmological scales.

Stellar streams. DM subhalos that cross stellar tidal streams—formed by the disruption of globular clusters or dwarf galaxies—can induce density gaps and spurs along the stream. The number and size of these features provide insight into the subhalo mass function [70,71]. Indeed, Refs. [72,73] find some of the tightest constraints on the WDM mass, deriving $m_{\text{WDM}} > 3.6$ keV (95% CL) by combining Gaia data with Pan-STARRS photometry of the GD-1 and Palomar 5 streams. Extracting DM subhalo properties from stream perturbations is, however, challenging due to degeneracies between the subhalo’s mass and its velocity [74,75]. Forthcoming data releases and improved

numerical simulations are expected to reduce these uncertainties and enhance our ability to extract subhalo properties to test DM scenarios [76–78].

UV luminosity functions. The UV luminosity function (UVLF), which measures the number density of galaxies as a function of UV magnitude, is a key probe of early galaxy formation. In WDM models, suppressed low-mass halo formation flattens the UVLF faint-end slope, enabling constraints from galaxy surveys at $z \gtrsim 4$ with HST and JWST [79–82]. Recent JWST data yield the strongest bound to date, $m_{\text{WDM}} \geq 3.2$ keV at 95% CL [83]. These constraints remain uncertain due to challenges in detecting high-redshift galaxies and modeling the halo mass–UV magnitude relation. Nevertheless, future JWST data will refine these limits and help disentangle baryonic effects [82,84].

Joint analyses. Finally, joint constraints can be derived by combining different small-scale probes. For instance, Ref. [85] combines posterior distributions from strong gravitational lensing, the Lyman- α forest, and MW satellite counts, obtaining $m_{\text{WDM}} > 6.048$ keV at 95% CL by marginalizing over nuisance parameters. Reference [86] combines posteriors from strong lensing and satellite data, finding a Bayesian lower bound of $m_{\text{WDM}} > 9.7$ keV at 95% CL. Joint analyses require careful treatment of nuisance parameters, as shared constraints must use consistent parametrizations or be properly disentangled to avoid bias. For our purposes, we therefore focus on individual constraints.

III. SETTING MASS BOUNDS ON QUASITHERMAL PSD: GENERAL STRATEGY

The primary goal of this work is to track the PSD of freeze-in DM and derive the corresponding mass bounds. Before turning to the details of the freeze-in mechanism in the next section, we outline the general strategy.

Regardless of the production mechanism, DM generation in the early universe eventually becomes inefficient, after which the evolution of the PSD is governed solely by cosmological redshift. We focus on this asymptotic form of the PSD and quantify its *warmness* through the second moment. To assess whether a given solution is viable, we compare its warmness to that of a thermal WDM particle whose mass saturates the observational bounds collected in Sec. II. The minimum mass for freeze-in DM is then defined as the value that yields the same warmness. This approach to setting mass bounds is justified by the thermal origin of the produced DM particles: although they never reach thermal equilibrium, they inherit typical energies comparable to the bath temperature.

A more comprehensive analysis would employ the full PSD to evaluate the DM impact on the formation of cosmological structures. However, this approach is computationally expensive for deriving bounds in concrete models and impractical for a systematic exploration of

their parameter space. Instead, we build on the results of Ref. [34], which established a connection between freeze-in DM and WDM. Using the Boltzmann solver CLASS [87], it was shown in a model-independent manner that the matter power spectrum for freeze-in DM exhibits features closely resembling those of WDM. In particular, the relative “transfer function”¹ displays a sharp cutoff at large wave numbers, indicating a suppression of small-scale structure that mirrors the behavior of WDM. As a result, the freeze-in transfer function for a given model can be effectively mapped onto that of WDM by selecting an appropriate value of m_{WDM} . Our strategy is thus supported qualitatively by the thermal origin of freeze-in DM and quantitatively by the results of Ref. [34].

In light of the above, freeze-in DM production is just one example where our strategy is effective. The same approach applies more broadly, provided two conditions are met. First, the DM PSD must be such that its physical effects are well captured by its first three moments.² Second, the first ($\langle p \rangle$) and second ($\langle p^2 \rangle$) moments must satisfy $\langle p \rangle^2 \sim \langle p^2 \rangle$. The freeze-in mechanism is a prototypical case where both conditions are naturally satisfied due to the thermal origin of the produced DM particles.

We refer to this class of DM candidates as having a *quasithermal* PSD. For illustration purposes, we provide an analytical expression for a generic DM candidate χ that captures its typical behavior

$$\text{quasithermal PSD: } f_{\chi}^{\text{qth}}(q) \propto \frac{q^{\alpha}}{e^{\beta q} + \kappa}. \quad (1)$$

First, we note that the quasithermal PSD is defined up to an overall normalization fixed by the relic density. Second, the Friedmann-Robertson-Walker (FRW) geometry ensures that the PSD depends only on FRW time and the magnitude of the physical momentum. We express the PSD in terms of the comoving momentum q , formally defined in the next paragraph. Importantly, q remains constant once DM production shuts off, naturally factoring out cosmological redshift. Consequently, the functional form in Eq. (1) describes the asymptotic, time-independent PSD in the late universe. The parameters (α, β, κ) are real, with κ restricted to $\{-1, 0, +1\}$. For instance, WDM corresponds to a thermal PSD with $(\alpha, \beta) = (0, 1)$ and κ determined by particle statistics. We require $\beta > 0$ for convergence of the normalization integral and $|\alpha| \lesssim 1$ to keep the PSD close to thermal. This is true for all freeze-in scenarios studied in this work. We stress once again that Eq. (1) is only for illustrative purposes and no results of this work depend on it.

¹Following standard conventions, we define the squared transfer function as the ratio of the linear power spectrum for freeze-in DM to that of Λ CDM: $T^2(k) \equiv P(k)/P_{\Lambda\text{CDM}}(k)$.

²The zeroth moment corresponds to the PSD normalization, fixed by the observed DM relic density.

We now return to the formal definition of comoving momentum, which will be essential in what follows. The spatial homogeneity and isotropy of the FRW geometry ensure that the PSD of the DM particle χ depends only on the FRW time t and the magnitude of the physical momentum $p(t)$. Even after interactions cease and Hubble expansion dominates, $p(t)$ remains time-dependent due to cosmic redshift, scaling as $p(t) \propto a(t)^{-1}$, where $a(t)$ is the FRW scale factor. To factor out this dependence, we introduce the dimensionless comoving momentum

$$q \equiv \frac{p(t)}{T_P} \frac{a(t)}{a(t_P)}. \quad (2)$$

Here, t_P is an *arbitrary* reference time, and T_P is the bath temperature at that time. While the choice of t_P is conventional and does not affect physical results, a suitable selection can simplify the analysis. In the freeze-in scenarios considered in this work, where production is IR dominated and proceeds through decays or scatterings mediated by renormalizable interactions, a convenient choice for T_P is the mass of the heaviest particle involved in the production process. Explicit examples of this choice are provided in the next section. As noted above, the PSD expressed in terms of the comoving momentum defined in Eq. (2) becomes time-independent once DM production has ceased.

The relation between t_P and T_P follows from the dynamics of the expansion. We assume a standard cosmological history where, at early times, the energy density is dominated by a hot plasma in thermal equilibrium. The scale factor evolves as $a(t) \propto t^{1/2}$, and the time-temperature relation is set by the Friedmann equation

$$H(t) = \frac{1}{2t} = \frac{\pi \sqrt{g_\star(T)}}{3\sqrt{10}} \frac{T^2}{M_{\text{pl}}}. \quad (3)$$

We adopt the results of Ref. [88] for the effective number of relativistic degrees of freedom, $g_\star(T)$. Likewise, the entropy density is $s = (2\pi^2/45)g_{\star s}(T)T^3$, where $g_{\star s}(T)$ is also taken from Ref. [88].

Even if DM particles never reach thermal equilibrium, it is convenient to define an effective DM temperature

$$T_\chi = T_P \frac{a(T_P)}{a(T)} = T \left(\frac{g_{\star s}(T)}{g_{\star s}(T_P)} \right)^{1/3}, \quad (4)$$

where $a(T)$ is the FRW scale factor expressed as a function of the bath temperature T via the relation in Eq. (3). The first equality allows us to recast the comoving momentum as $q = p(t)/T_\chi$, while the second follows from entropy conservation ($sa^3 = \text{const}$). As pointed out by Refs. [25–28,31] and quantitatively confirmed by Ref. [34], a key quantity for cosmological structure formation is the DM root-mean-square velocity

$$W_\chi \equiv \frac{\sqrt{\langle p^2 \rangle}}{m_\chi} = \frac{T_\chi}{m_\chi} \sigma_q, \quad (5)$$

where the dimensionless PSD second moment is

$$\sigma_q \equiv \left(\frac{\int dq q^4 f_\chi(q)}{\int dq q^2 f_\chi(q)} \right)^{1/2}. \quad (6)$$

Although T_χ and σ_q both depend explicitly on the arbitrary choice of T_P , physical results remain unaffected. This is evident from Eq. (5): their product corresponding to the r.m.s. physical momentum is invariant. As we will make explicit in the next paragraph, W_χ is the key quantity for setting mass bounds, confirming that our results are independent of the choice of T_P .

Having established the necessary framework, we now determine DM mass bounds by comparing the warmness W_χ of a given model to the corresponding W_{WDM} for the WDM scenario. For convenience, all relevant WDM results are summarized in Appendix A, including the PSD shape, the relic density constraint that fixes the WDM temperature, and the calculation of W_{WDM} as defined in Eq. (5). The mass bound is obtained as follows. We start from a lower bound $m_{\text{WDM}}^{\text{min}}$ taken from the list reported in Sec. II. This minimum WDM mass corresponds to a maximum warmness $W_{\text{WDM}}^{\text{max}}$. A quasithermal DM candidate must then satisfy $W_\chi \leq W_{\text{WDM}}^{\text{max}}$, yielding the minimum allowed DM mass

$$m_\chi^{\text{min}} = 22 \text{ keV} \left(\frac{m_{\text{WDM}}^{\text{min}}}{6 \text{ keV}} \right)^{4/3} \left(\frac{\sigma_q}{3.6} \right) \left(\frac{106.75}{g_{\star s}(T_P)} \right)^{1/3}. \quad (7)$$

Rescaling a given WDM mass bound $m_{\text{WDM}}^{\text{min}}$ therefore requires knowledge of two parameters: the dimensionless second moment of the PSD, σ_q , and the temperature introduced to define the comoving momentum, T_P . This might suggest that the resulting mass bound depends on the (arbitrary) choice of T_P . However, as already emphasized, physical results cannot depend on this choice. Indeed, the dimensionless second moment σ_q depends on the choice of T_P , and this dependence exactly cancels out with that of the $g_{\star s}(T_P)$ factor, rendering the mass bound in Eq. (7) independent of T_P .

We summarize our results in the general form shown in Fig. 3. Taking $m_{\text{WDM}}^{\text{min}} = 6.8 \text{ keV}$ as a reference value, we display isocontours of m_χ^{min} in the (σ_q, T_P) plane. The interpretation of this plane is as follows: given a specific DM model with fixed masses and couplings, one selects a convenient temperature T_P and solves the Boltzmann equation to obtain the PSD. This allows for the computation of σ_q , identifying a point in the plane. Repeating the procedure with a different T_P leads to a different value of σ_q , thereby shifting the model's location in the plane. In summary, each model traces out a line in the (σ_q, T_P) plane




FIG. 3. Mass bounds on DM with a quasithermal PSD in the (σ_q, T_P) plane. The dimensionless second moment σ_q is defined in Eq. (6), while the temperature T_P appears in the definition of the comoving momentum, Eq. (2). Contours show the minimum allowed DM mass, obtained via the rescaling in Eq. (7) using $m_{\text{WDM}}^{\text{min}} = 6.8$ keV as the reference value.

that follows the shape of the isocontours in Fig. 3, which represent constant values of W_χ .

Associating T_P with the bath temperature at which DM production is most efficient clarifies the behavior of the isocontours in Fig. 3. The weakest mass bounds appear in the upper-left region of the plot. This is expected: moving left along the horizontal axis corresponds to smaller values of σ_q , indicating colder DM distributions with lower velocity dispersion. The vertical dependence is similarly intuitive: larger values of T_P correspond to earlier production times, resulting in more redshifting and cooling of the DM particles, especially when the thermal bath crosses particle mass thresholds. Notably, the bound becomes significantly more stringent near the GeV scale, where a sharp drop in the number of SM degrees of freedom enhances cooling and leads to a smaller σ_q .

IV. FREEZE-IN PRODUCTION

After the general discussion about DM candidates with a quasithermal PSD, we now turn to the specific production mechanism known as freeze-in. In this scenario, the early Universe is assumed to be reheated after inflation with a negligible abundance of DM. Particles in the thermal bath can occasionally scatter or decay into final states that include DM. As long as the relevant interactions are sufficiently feeble, the inverse (destruction) processes remain negligible and DM never reaches thermal equilibrium. The resulting DM population simply redshifts with the expansion of the universe, propagating along geodesics and persisting to the present time.

We follow the notation and conventions of Ref. [34], and summarize here only the key results. The general freeze-in

production process, involving n bath particles in the initial state and a final state composed of m bath particles and ℓ DM particles, reads

$$\underbrace{\mathcal{B}_1 + \dots + \mathcal{B}_n}_n \rightarrow \underbrace{\mathcal{B}_{n+1} + \dots + \mathcal{B}_{n+m}}_m + \underbrace{\chi + \dots + \chi}_\ell. \quad (8)$$

Bath particles $\mathcal{B}_i$ are described by equilibrium distributions $f_{\mathcal{B}_i}^{\text{eq}}(k_i) = [\exp(E_i/T) \pm 1]^{-1}$, either characterized by Fermi-Dirac (+) or Bose-Einstein (−) statistics. The PSD evolution is governed by the Boltzmann equation

$$g_\chi \frac{df_\chi(p, t)}{dt} = g_\chi \frac{\sum_\alpha C_\alpha(t)}{E}. \quad (9)$$

The sum over α runs over all contributing production processes. For freeze-in, the right-hand side of the Boltzmann equation is a function of the bath temperature T and not a functional of the PSD f_χ . Its explicit form for the process in Eq. (8) is

$$C_{n \rightarrow m\ell} \equiv \frac{\ell}{2} \int \prod_{i=1}^{n+m} d\mathcal{K}_i \prod_{j=2}^{\ell} d\Pi_j (2\pi)^4 \delta^{(4)}(P_f - P_i) |\mathcal{M}_\alpha|^2 \times \prod_{i=1}^n f_{\mathcal{B}_i}(k_i) \prod_{j=n+1}^{n+m} [1 \mp f_{\mathcal{B}_j}(k_j)]. \quad (10)$$

Here, P_i and P_f denote the total four-momenta of the initial and final states, respectively. The measures $d\mathcal{K}_i$ and $d\Pi_j$ are the Lorentz-invariant phase space elements, given by $g_i d^3 p_i / [(2\pi)^3 2E_i]$, for bath and DM particles, respectively. The quantity $|\mathcal{M}_\alpha|^2$ is the squared matrix element for the process, averaged over both initial and final state degrees of freedom. This, along with the number and type of production channels, constitutes the only model-dependent input for freeze-in.

Before solving the Boltzmann equation, it is convenient to change the time and momentum variables. First, we replace the cosmic time t with the thermal bath temperature T ; the corresponding Jacobian factor is determined from entropy conservation. For the momentum, we adopt the comoving variable introduced in Eq. (2), which naturally factors out the effects of cosmic redshift.

The Boltzmann equation is integrated from an initial temperature T_{UV} , chosen to be much larger than the reference temperature T_P at which production is most efficient, down to the present temperature T_0 . This yields the asymptotic comoving distribution

$$g_\chi f_\chi = \int_{T_{\text{UV}}}^{T_0} \frac{d \log T}{H(T)} \left(1 + \frac{1}{3} \frac{d \log g_{\star s}}{d \log T} \right) g_\chi \frac{\sum_\alpha C_\alpha}{E}. \quad (11)$$

We impose the initial condition $f_\chi(T_{\text{UV}}) = 0$ and focus on scenarios in which the final relic abundance is insensitive to

the precise value of T_{UV} . It has been shown that this IR dominated production happens always for decays and for scattering mediated by renormalizable interactions [6–9]. We safely take the $T_{\text{UV}} \rightarrow \infty$ limit in the integral above. Moreover, if we are interested in the asymptotic form of the PSD expressed in terms of the comoving momentum, there is no need to integrate all the way down to the present temperature T_0 . Upon identifying T_P with the mass of the heaviest particle involved in the production process, it is sufficient to stop the integration at a temperature well below T_P , where the Maxwell-Boltzmann exponential suppression of the bath particle density renders further production negligible.

Whether DM production is single (i.e., $\ell = 1$) or multiple (i.e., $\ell > 1$), the functional shape of the PSD resulting from the integral in Eq. (11) is not substantially affected. The main difference is the multiplicative factor of ℓ in the collision term, Eq. (10), which is reabsorbed by the normalization imposed by the relic density constraint. Minor corrections may arise from the phase space integrals, as multiple production requires integration over the momenta of the additional DM particles. This paper focuses on the regime where DM is much lighter than the heaviest bath particle involved in production. In this regime, and absent nearly degenerate bath particles, integrating over additional DM momenta is equivalent to integrating over momenta of light bath particles. While the detailed functional form may differ, in the

IR-dominated regime the matrix element can be taken as constant to excellent approximation. Pauli-blocking or Bose-Einstein effects for final-state bath particles introduce corrections to the second moment σ_q at the percent level [34]. For these reasons, we concentrate here on single production, while multiple production is treated in the next section in the context of explicit microscopic models.

In what follows, we analyze three distinct freeze-in production channels: two-body decays, binary scatterings, and three-body decays. The analysis in this section is model-independent and depends on a single parameter: the overall freeze-in mass scale, set by the heaviest bath particle involved in the production process. Without loss of generality, we denote this particle as $\mathcal{B}_1$, with mass m_1 , thereby defining the production temperature as $T_P = m_1$. The interaction strength is parametrized through a constant matrix element, whose value is fixed by the relic density constraint. All particle statistics are taken to follow Maxwell-Boltzmann distributions. For each production channel, we solve the Boltzmann equation and extract the corresponding value of σ_q .

Our results are presented in Fig. 4. For each case, the masses of the other bath particles are chosen such that m_2/m_1 is set by the legend, and $m_3 = 0$. The background (black) contours in all three panels reproduce those of Fig. 3, which show DM mass bounds in the (σ_q, T_P) plane with the only difference that here the vertical axis is relabeled as m_1 . A specific choice for the mass spectrum




FIG. 4. Lower mass bounds on freeze-in produced DM from: two-body decays (left), binary scatterings (center), and three-body decays (right). The (σ_q, m_1) plane is the same as in Fig. 3, with T_P replaced by the mass of the heaviest particle involved in production, m_1 . Solid black contours reproduce those of Fig. 3, corresponding to the reference WDM bound $m_{\text{WDM}}^{\text{min}} = 6.8$ keV. For each channel, we explore different mass spectra, as indicated in the legend, by varying $m_2/m_1 = \{0, 0.25, 0.5, 0.75, 0.9\}$. Whenever a third bath particle is involved, we always set $m_3 = 0$. Each choice for the mass spectrum identifies a point in the plane; the black contour intersecting the point gives the associated mass bound. Vertical dashed lines show analytically computed values of σ_q (see Appendix B), and the thick gray line marks the WDM benchmark $\sigma_q^{\text{WDM}} \simeq 3.6$.




FIG. 5. Lower bounds on the DM mass $m_\chi^{\min}$ for the three production mechanisms considered in this work: two body decays (left), binary scatterings (center), and three body decays (right). The three lines in each panel correspond to the reference WDM mass bounds shown in Fig. 2, using the same color scheme. Matrix elements are taken to be constant and adjusted to reproduce the observed DM relic abundance. The variable m_1 denotes the mass of the heaviest particle participating in the process, while B_2 (and B_3 when present) are assumed to be massless. The DM mass bound is shown as a function of m_1 .

corresponds to a colored point in this plane. In the evaluation of the collision terms, the DM mass is chosen to satisfy $m_\chi \ll m_1$ and is taken to lie slightly above the corresponding bound.

(i) *Two-body decays.* The case of two body decays $B_1 \rightarrow B_2 \chi$ is the only one in which the constant matrix element approximation holds exactly. We observe that, except in the region where m_1 lies between 0.1 and 1 GeV and near the MeV scale, the bounds on m_χ are less stringent compared to the thermal scenario, indicated by the vertical thick gray line. Red points correspond to $m_2 = 0$, while purple points indicate $m_2 = 0.9m_1$, which reduces the energy available for DM particles. This reduction, referred to as *phase space suppression*, weakens the warmness constraints as expected.

(ii) *Binary scattering.* In this case, we consider DM production through three scattering processes ($B_1 B_2 \rightarrow B_3 \chi$ and its permutations). All matrix elements are approximated as constants. We present results in the (σ_q, m_1) plane for different choices of the mass spectrum once the largest scale m_1 is fixed. The color coding corresponds to different values of the mass ratio m_2/m_1 . The results correctly reflect that phase space suppression cannot occur in this case, since we include all permutations of the initial and final state bath particles. The bounds on m_χ are again weaker compared to the thermal PSD case. Notably, we find that the mass bounds depend only mildly on m_2 and m_3 , with the dominant factor being m_1 , which sets the typical production energy. For m_1 greater than a GeV, we find $m_\chi^{\min} \simeq 17$ keV, while for m_1 below 100 MeV, we obtain $m_\chi^{\min} \simeq 30$ keV.

(iii) *Three-body decays.* The case of three body decays $B_1 \rightarrow B_2 B_3 \chi$ shares features with both scenarios discussed above. We again approximate the matrix elements as constants. Since this is a decay, the total energy in the final state is constrained by the mass of the decaying particle. As a result, warmness bounds can be relaxed by phase space suppression, achieved by choosing m_2 sufficiently close to m_1 . However, even when m_2 is large, the resulting bounds are generally weaker than those obtained for two-body decays. Specifically, for m_1 greater than a GeV, we find $m_\chi^{\min} \simeq 13$ keV, while for m_1 below 100 MeV, we obtain $m_\chi^{\min} \simeq 27$ keV.

Finally, we derive lower bounds on the DM mass m_χ using Eq. (7), taking into account different values for the WDM mass bound $m_{\text{WDM}}^{\min}$. The resulting constraints are displayed in Fig. 5 for the three production mechanisms discussed above, with m_1 varied and all other particles taken to be massless (except the DM). Among these scenarios, two-body decays yield the strongest bounds, followed by binary scatterings and then three-body decays.

V. MINIMAL TWO-BODY DECAYS

After the model independent analysis of freeze-in presented in the previous section, we now derive mass bounds for explicit examples of microscopic theories. In this section, we focus on two *minimal* scenarios, each based on a renormalizable extension of the SM, that offer a viable DM candidate produced via two-body decays. It is important to recall that the results of the previous section were obtained under two simplifying assumptions: constant matrix elements and Maxwell-Boltzmann statistics for all

particles involved. The first assumption is well justified in the case of two-body decays, as the final state is monochromatic. Once a specific microscopic model is chosen, however, we can go beyond these approximations by solving the Boltzmann equation with the exact collision operators and including all quantum statistical effects. As shown in Ref. [34], these quantum corrections induce only a small change in the parameter σ_q defined in Eq. (6), typically at the percent level. Consequently, the resulting shift in the mass bound derived from Eq. (7) is negligible. This confirms that the model independent results obtained previously remain reliable even in the context of complete microscopic models.

A. Scalar DM via the Higgs portal

The first case we consider is the most minimal extension of the SM in terms of new degrees of freedom. We introduce one real scalar field ϕ , neutral under all SM gauge symmetries. To ensure the stability of ϕ , we impose a discrete $\mathbb{Z}_2$ symmetry under which all SM fields are even and ϕ is odd. The most general set of renormalizable interactions involving ϕ that can be added to the SM Lagrangian is given by

$$\mathcal{L}_\phi = \frac{1}{2}(\partial_\mu \phi)^2 - \frac{1}{2}m_\phi^2 \phi^2 - \frac{\lambda_{\text{mix}}}{2} H^\dagger H \phi^2 - \frac{\lambda_\phi}{4!} \phi^4. \quad (12)$$

The first two terms describe a canonically normalized real scalar with mass m_ϕ . To preserve the $\mathbb{Z}_2$ symmetry and guarantee the stability of ϕ , we must prevent it from acquiring a vacuum expectation value (vev). In the regime $\lambda_{\text{mix}} \ll 1$, which is required for freeze-in production, the scalar potential can be analyzed by decoupling the dynamics of the Higgs doublet. In this limit, the condition for a vanishing vacuum expectation value is simply $\lambda_\phi > 0$ and $m_\phi^2 > 0$.

DM production in the early universe is controlled by the interaction between the field ϕ and the SM Higgs boson h , identified as $H = (0(v_H + h)/\sqrt{2})^T$, where the Higgs vev is determined from the Fermi constant and evaluates to $v_H \simeq 246$ GeV. The relevant production channels for DM are mediated by the cubic interaction term $h\phi^2$, which emerges from expanding the scalar potential around the Higgs vev. This setup has been shown to reproduce the observed DM abundance via both freeze-out [89,90] and freeze-in [6]. The freeze-out regime is tightly constrained by direct detection experiments.

For DM masses below half the Higgs mass, the decay $h \rightarrow \phi\phi$ is kinematically allowed, with decay width

$$\Gamma_{h \rightarrow \phi\phi} = \frac{\lambda_{\text{mix}}^2 v_H^2}{32\pi m_h} \sqrt{1 - \frac{4m_\phi^2}{m_h^2}}. \quad (13)$$

Since we focus on freeze-in production of DM in the tens of keV range, this decay channel is always open. Moreover,

it has been shown that the contribution from scatterings is subdominant and can be safely neglected when Higgs decays are kinematically accessible [91,92]. This setup thus provides a minimal and concrete scenario in which DM is produced via freeze-in through the two-body decay of a SM particle with known mass.

The squared matrix element for a two-body decay is constant and directly proportional to the decay width. We fix its value (i.e., the coupling λ_{mix}) to reproduce the observed DM relic density, leaving the DM mass as the only free parameter subject to constraints. The lower bound on m_ϕ can be inferred from the left panel of Fig. 5, selecting $m_1 = m_h \simeq 125$ GeV, with the caveat that those results were obtained under the assumption of Maxwell-Boltzmann statistics and constant matrix elements. However, quantum statistical corrections affect the mass bound only at the percent level, making the Maxwell-Boltzmann approximation sufficiently accurate. The resulting bounds on the DM mass are

$$m_\phi > \begin{cases} 23 \text{ keV} & \text{MW satellites} \\ 20 \text{ keV} & \text{JWST lensing} \\ 18 \text{ keV} & \text{Lyman } \alpha \end{cases}. \quad (14)$$

B. Sterile neutrinos

The second model we study is also guided by the criterion of minimality, but now we introduce a new fermionic degree of freedom in the form of a Weyl field N . As in the previous case, the new field is assumed to be neutral under all SM gauge symmetries. This choice significantly complicates the construction of a viable Lagrangian if we continue to restrict ourselves to renormalizable theories. The most general set of renormalizable terms that can be added to the SM Lagrangian reads

$$\mathcal{L}_N = N^\dagger i \bar{\sigma}^\mu \partial_\mu N - \left[\frac{m_N}{2} N N + \lambda_\nu L H N + \text{H.c.} \right]. \quad (15)$$

In addition to the canonical kinetic term, there are two contributions: a Majorana mass term, which is allowed since N is a complete SM singlet, and a Yukawa interaction involving the SM gauge-invariant combination LH , where L denotes the lepton doublet. We omit flavor indices in the Lagrangian above. In principle, one could consider scenarios in which N couples to all SM lepton generations, promoting the Yukawa coupling to a vector, or even introduce multiple fermion singlets, resulting in a full Yukawa matrix. However, the conclusions of this subsection are not sensitive to these model-building choices.

The scenario illustrated by the above Lagrangian is precisely what is required to generate masses for the SM *active neutrinos* through the addition of fermion singlets. Moreover, if the *sterile neutrino* field N has a mass in the tens of keV range, it remains stable on cosmological

timescales and can serve as a viable DM candidate. An unavoidable production mechanism for sterile neutrinos in the early universe is the Dodelson-Widrow process [93], in which oscillations of active neutrino fields ν populate the sterile sector. However, the scenario in which sterile neutrino DM is produced solely via this mechanism is excluded by x-ray observations, which constrain the radiative decay channel $N \rightarrow \gamma\nu$ by searching for monochromatic photon lines [94,95]. A possible way to evade these bounds is to invoke an alternative DM genesis via resonant production through the Shi-Fuller mechanism [96] or nonstandard neutrino self-interactions [97,98].

In this work, we consider another production channel. The Yukawa interaction in Eq. (15) can be made invariant under lepton number if the field N is appropriately charged. However, the Majorana mass term violates lepton number by two units. This low-energy theory can arise from a lepton-number-conserving ultraviolet completion by introducing a new complex scalar Φ , which is also a singlet under the SM gauge group. The Yukawa operator ΦNN conserves lepton number provided Φ carries the appropriate charge. If Φ acquires a nonvanishing vev, it generates the Majorana mass term for N at low energies. Expanding the scalar field around its vev as $\Phi = (v_\Phi + \rho)/\sqrt{2}$, we identify the radial mode ρ , which can thermalize with the SM bath through a quartic interaction with the Higgs doublet H . The decay of this thermalized ρ field into sterile neutrinos then acts as a freeze-in production mechanism. We focus on the regime where this is the dominant channel for sterile neutrino DM production [99]. Unlike the minimal Higgs portal scenario discussed earlier, here the mass of the decaying particle m_ρ remains a free parameter. We derive the following lower bounds on the DM mass

$$m_N > \begin{cases} \begin{cases} 47 \text{ keV} & \text{MW Satellites} \\ 40 \text{ keV} & \text{JWST} \\ 37 \text{ keV} & \text{Lyman-}\alpha \end{cases} & m_\rho = 100 \text{ MeV} \\ \begin{cases} 24 \text{ keV} & \text{MW Satellites} \\ 21 \text{ keV} & \text{JWST} \\ 19 \text{ keV} & \text{Lyman-}\alpha \end{cases} & m_\rho = 10 \text{ GeV} \\ \begin{cases} 22 \text{ keV} & \text{MW Satellites} \\ 19 \text{ keV} & \text{JWST} \\ 17 \text{ keV} & \text{Lyman-}\alpha \end{cases} & m_\rho = 1 \text{ TeV} \end{cases} \quad (16)$$

As in the previous scenario, we verified that including the Bose-Einstein statistical factors for the decaying scalar alters the mass bounds only at the percent level.

VI. ALP COUPLED TO SM FERMIONS

We now consider scenarios where DM production occurs via both two-body decays and binary scatterings. The DM

candidate is a CP -odd scalar a , originating from the spontaneous breaking of a global symmetry. A prominent example is the QCD axion associated with Peccei-Quinn symmetry breaking [100–103]. Adopting a low-energy effective field theory, we impose a shift symmetry consistent with the Nambu-Goldstone nature of a , and focus on couplings of this ALP to SM fermions only

$$\mathcal{L}_a = \frac{1}{2}(\partial_\mu a)^2 - \frac{1}{2}m_a^2 a^2 + \mathcal{L}_a^{(\text{FC})} + \mathcal{L}_a^{(\text{FV})}. \quad (17)$$

Here, $\mathcal{L}_a^{(\text{FC})}$ and $\mathcal{L}_a^{(\text{FV})}$ contain flavor-conserving and violating interactions, respectively. The mass term softly breaks the shift symmetry and is a free parameter.

A. Flavor-conserving couplings

We consider the flavor-conserving interactions defined in the basis where the SM Yukawa matrices are diagonal

$$\mathcal{L}_a^{(\text{FC})} = -\frac{\partial_\mu a}{2\Lambda} \sum_\psi C_\psi \bar{\psi} \gamma^\mu \gamma^5 \psi. \quad (18)$$

It is convenient to define the combination $\Lambda_\psi \equiv \Lambda/C_\psi$.

For ALP coupled to leptons, $\psi = \ell$, the production channels are lepton/antilepton annihilations $\ell^+ \ell^- \rightarrow \gamma a$ and Compton-like scattering $\ell^\pm \gamma \rightarrow \ell^\pm a$. The squared matrix elements expressed in terms of the Mandelstam variables result in [104–106]

$$|\mathcal{M}_{\ell^+ \ell^- \rightarrow \gamma a}|^2 = \frac{e^2}{2\Lambda_\ell^2} \frac{m_\ell^2 s^2}{(m_\ell^2 - t)(s + t - m_\ell^2)}, \quad (19a)$$

$$|\mathcal{M}_{\ell^\pm \gamma \rightarrow \ell^\pm a}|^2 = \frac{e^2}{2\Lambda_\ell^2} \frac{m_\ell^2 t^2}{(s - m_\ell^2)(s + t - m_\ell^2)}, \quad (19b)$$

where e is the electron charge.

The analysis for ALPs coupled to quarks proceeds analogously, with the key difference that we restrict to the three heaviest quarks to avoid issues due to confinement. The relevant squared matrix elements are

$$|\mathcal{M}_{q\bar{q} \rightarrow \gamma a}|^2 = \frac{g_s^2}{36\Lambda_q^2} \frac{m_q^2 s^2}{(m_q^2 - t)(s + t - m_q^2)}, \quad (20a)$$

$$|\mathcal{M}_{qg \rightarrow qa}|^2 = \frac{g_s^2}{36\Lambda_q^2} \frac{m_q^2 t^2}{(s - m_q^2)(s + t - m_q^2)}, \quad (20b)$$

where g_s denotes the strong coupling constant. For the top quark, whose mass lies well above the QCD scale, freeze-in calculations are fully under control. For the bottom and charm quarks, we conservatively stop the Boltzmann evolution at $T_{\text{stop}} = 1 \text{ GeV}$ and neglect production at lower temperatures. Similar issues have been discussed in the context of axion dark radiation [106,107], where thermal

TABLE I. Summary of the results for the ALP DM scenario in Eq. (17) with flavor-conserving couplings to SM fermions. Freeze-in production takes place via binary scatterings. We consider both couplings to charged leptons (first three rows) and to heavy quarks (last three rows). For each choice of the SM fermion ψ , we find the mass bounds $m_a^{\min}$ from three different observables and report their values. For ALP masses in this ballpark range, the DM relic density constraint fixes the value of the axion coupling to $\Lambda_{\psi}^{\text{relic}}$. We also determine the radiatively induced coupling $g_a^{1\text{-loop}}$ to photons and the resulting ALP lifetime $\tau_{a \rightarrow \gamma\gamma}$. Finally, in the last column we state whether this scenario is allowed by x-ray constraints.

ψ	$m_a^{\min}$ (keV)			Relic density requirement and x-ray constraints			
	Ly α	JWST	Satellite	$\Lambda_{\psi}^{\text{relic}}$ (GeV)	$-\Delta g_{a\gamma\gamma}$ (GeV $^{-1}$)	$\tau_{a \rightarrow \gamma\gamma}$ (sec)	Allowed?
Electron	39	43	49	5.7×10^7	3.1×10^{-14}	1.2×10^{18}	$\times$
Muon	33	37	43	3.6×10^8	5.6×10^{-20}	1.1×10^{30}	$\checkmark$
Tau	19	21	24	4.9×10^8	7.3×10^{-23}	1.8×10^{36}	$\checkmark$
Charm	19	20	23	4.5×10^9	1.8×10^{-23}	3.1×10^{37}	$\checkmark$
Bottom	19	21	24	7.3×10^9	3.2×10^{-25}	8.3×10^{40}	$\checkmark$
Top	18	19	22	2.8×10^{10}	1.5×10^{-28}	5.5×10^{47}	$\checkmark$

axions lead to potentially observable contributions to ΔN_{eff} . Here, by contrast, we are interested only in the *shape* of the resulting PSD, which depends on the temperature evolution of each individual momentum bin in the collision term. Reference [108] showed how this dependence is mild, supporting the reliability of our results.

ALP-fermion interactions induce one-loop couplings to photons, rendering the ALP unstable and enabling observational tests through x-ray line searches. The effective interaction generated at one loop is given by [109,110]

$$\Delta \mathcal{L}_a^{(1\text{-loop})} = -\frac{\Delta g_{a\gamma\gamma}}{4} a F^{\mu\nu} \tilde{F}_{\mu\nu}, \quad (21a)$$

$$\Delta g_{a\gamma\gamma} \simeq -\frac{\alpha_{\text{em}} N_c^{\psi} Q_{\psi}^2}{12\pi \Lambda_{\psi}} \left(\frac{m_a}{m_{\psi}} \right)^2. \quad (21b)$$

This leads to a decay width $\Gamma_{a \rightarrow \gamma\gamma} = \Delta g_{a\gamma\gamma}^2 m_a^3 / (64\pi)$. We summarize the resulting constraints for ALP DM produced via flavor-conserving scatterings with SM fermions in Table I. Notably, the case of ALPs coupled to electrons is excluded by x-ray observations [111–114].

B. Flavor-violating couplings

We now consider the case in which the ALP couplings are flavor-violating with interactions

$$\mathcal{L}_a^{(\text{FV})} = -\frac{\partial_{\mu} a}{2\Lambda} \sum_{\psi \neq \psi'} \bar{\psi}' \gamma^{\mu} (\mathcal{V}_{\psi\psi'} + \mathcal{A}_{\psi\psi'} \gamma^5) \psi + \text{H.c.}, \quad (22)$$

where $\mathcal{V}_{\psi\psi'}$ and $\mathcal{A}_{\psi\psi'}$ encode the vector and axial-vector couplings, respectively. These interactions have been considered to exploit the consequences of flavor-violating couplings not only for ALP DM [115,116] but also for ALP dark radiation [117]. We define the effective scale

$$\Lambda_{\psi\psi'} \equiv \frac{\Lambda}{\sqrt{|\mathcal{V}_{\psi\psi'}|^2 + |\mathcal{A}_{\psi\psi'}|^2}}. \quad (23)$$

The squared matrix element for $\psi \rightarrow \psi' a$ reads [104,107]

$$|\mathcal{M}_{\psi \rightarrow \psi' a}|^2 = \frac{m_{\psi}^4}{8N_c^{\psi} \Lambda_{\psi\psi'}^2} \left(1 - \frac{m_{\psi'}^2}{m_{\psi}^2} \right)^2, \quad (24)$$

with N_c^{ψ} the number of colors of the fermion ψ .

We present the results for ALP production via flavor-violating decays of SM fermions in Table II. The cases involving charged leptons pose no complications from strong interactions and yield reliable mass bounds, fully consistent with our model-independent analysis. Quark decays, on the other hand, are more delicate due to confinement effects. The result for the top quark, owing to its large mass, remains robust and in agreement with the general analysis. The situation differs significantly for

TABLE II. Results for the ALP DM scenario described by Eq. (22), with flavor-violating couplings to SM fermions. Freeze-in production occurs via two-body decays $\psi \rightarrow \psi' a$. We consider couplings to charged leptons and the top quark. Decays of bottom and charm quarks are omitted for reasons discussed in the main text. For each SM fermion ψ , we report the lower bound on the ALP mass $m_a^{\min}$ from three different observables, along with the value of the coupling scale $\Lambda_{\psi\psi'}^{\text{relic}}$ defined in Eq. (23) required to reproduce the relic abundance.

ψ	$m_a^{\min}$ (keV)			$\Lambda_{\psi\psi'}^{\text{relic}}$ (GeV)
	Ly α	JWST	Satellite	
Tau	25	27	31	2.9×10^9
Muon	37	40	47	2.9×10^9
Top	18	20	22	3.4×10^{10}

bottom and charm quarks. As in the case of freeze-in production via scatterings with flavor-conserving couplings, we adopt a conservative approach by halting the Boltzmann integration at $T_{\text{stop}} = 1$ GeV to avoid uncertainties near the QCD phase transition. This leads to surprisingly weaker mass bounds, especially for the charm quark, compared to the model-independent expectations. However, this discrepancy is not unexpected. As shown in Ref. [108], the collision term's temperature dependence at fixed momentum is quite pronounced for decays, unlike in the scattering case. In particular, high-momentum modes (relevant at lower temperatures) are artificially removed by the choice $T_{\text{stop}} = 1$ GeV, resulting in a colder and narrower PSD, which in turn relaxes the mass bounds. These results are not reliable, as they are sensitive to the detailed *shape* of the PSD, which cannot be accurately determined when the integration is artificially truncated. For this reason, we choose to omit the bottom and charm quark cases from Table II.

VII. DARK PHOTON PORTAL

The last theory we consider features a SM-singlet Dirac fermion χ as the DM candidate, which interacts with the visible sector through a new spin-one boson A'_μ . The interaction arises via kinetic mixing between the Abelian field strengths of the dark photon (DP) mediator, $F'_{\mu\nu}$, and the SM photon, $F_{\mu\nu}$ [22,32,118,119]. The Lagrangian relevant for the phenomenology of this framework is

$$\begin{aligned} \mathcal{L}_{\text{DP}} = & -\frac{1}{4}F'_{\mu\nu}F'^{\mu\nu} + \frac{\kappa}{2}F'_{\mu\nu}F^{\mu\nu} + \frac{1}{2}m_{A'}^2 A'_\mu A'^\mu \\ & + eJ_{\text{EM}}^\mu A_\mu + \bar{\chi}\gamma^\mu (i\partial_\mu + e_D A'_\mu)\chi. \end{aligned} \quad (25)$$

Here, A_μ denotes the SM photon, which couples to the electromagnetic current J_{EM}^μ composed solely of SM fields, since χ is a gauge singlet. The Abelian gauge couplings of the visible and dark sectors are denoted by e and e_D , respectively.

After transforming to a canonical basis, the DM field acquires a millicharge $Q_\chi = \kappa e_D/e$. This induces an effective interaction that mediates the DM pair-production process $e^+e^- \rightarrow \bar{\chi}\chi$, where we focus on electron-positron annihilations, as heavier fermions give negligible contributions. The squared matrix element, averaged over initial and final spins, is given by

$$\begin{aligned} |\mathcal{M}_{e^+e^- \rightarrow \bar{\chi}\chi}|^2 = & \frac{Q_\chi^2 e^4}{2s^2} (s^2 + 2st - t(m_e^2 + m_\chi^2) \\ & + 2t^2 + 2(m_e^2 + m_\chi^2)^2). \end{aligned} \quad (26)$$

A competing process in DM production is plasmon decay. In the thermal plasma, the photon acquires an effective mass from thermal corrections, allowing the decay $\gamma^* \rightarrow \bar{\chi}\chi$ to occur. The calculation of the corresponding production

rate requires special care, and we adopt the results of Ref. [22].

Once both production channels are taken into account, the coupling Q_χ can be fixed by requiring that χ reproduces the observed relic abundance. Interestingly, this model allows for DM production via freeze-in through two distinct and competing processes with very different characteristics. While the number density is always dominated by one of the two, making the other a subleading correction, the impact on the phase-space distribution is more subtle. Accounting for both e^+e^- annihilations and plasmon decays is essential. The PSD resulting from plasmon decays is very narrow and peaked at low comoving momenta, whereas the distribution produced by e^+e^- annihilations is much broader and characterized by a higher average momentum. As a result, even a small contribution from annihilations—particularly relevant in the low- m_χ regime—can significantly “warm up” the PSD, affecting the bound on m_χ at the level of 50%. Another noteworthy feature of this model is that the production scale depends explicitly on the DM mass. Since plasmon decays occur when the plasma frequency satisfies $\omega_p(T) \simeq 2m_\chi$, the production temperature is approximately $T \approx 20m_\chi$. This behavior is somewhat unconventional for DM produced via IR freeze-in.

Figure 6 illustrates how to identify the DM mass bound. The upper panel displays the relative contributions to the relic density from DM produced via electron-positron annihilation (red) and plasmon decay (blue), as a function of the DM mass. The lower panel shows the ratio between the r.m.s. velocity of the χ particles and that of the reference WDM model. We show the individual contributions as well as the result obtained from the full PSD. The mass bound is obtained by identifying the intersection point where $W_\chi/W_{\text{WDM}} = 1$, and we find

$$m_\chi > \begin{cases} 31 \text{ keV} & \text{MW satellites} \\ 26 \text{ keV} & \text{JWST lensing} \\ 24 \text{ keV} & \text{Lyman } \alpha \end{cases}. \quad (27)$$

We conclude this section with a comparison from the analysis of Ref. [32] on the same model. Our bounds are roughly a factor of two more stringent, and we identify this discrepancy as due to different implementations of the criterion used to determine the minimum value of m_χ consistent with WDM clustering bounds. In Ref. [32], m_χ is chosen such that the half-mode (defined as the scale where the DM power spectrum is suppressed by 50% relative to CDM) matches that of the WDM spectrum. This methodology is supported by what highlighted by Ref. [120]. However, as evident in Fig. 3 of Ref. [32], the resulting allowed power spectrum for the nonthermal distribution contains a high-power region that lies below the




FIG. 6. Lower mass bounds for the dark photon portal. The bottom panel shows the ratio W_χ/W_{WDM} as a function of the DM mass, highlighting the individual contributions from plasmon decay and e^+e^- annihilation. The upper panel displays the corresponding contributions to the relic density. While the relic abundance is predominantly set by plasmon decays, the lower mass bound $m_\chi^{\min} = 31.3$ keV (obtained via the condition $W_\chi/W_{\text{WDM}} < 1$ and identified by the vertical gray dashed line) lies between the regions where the two production mechanisms dominate. The reference WDM constraint used is $m_{\text{WDM}}^{\min} = 6.8$ keV.

corresponding WDM spectrum and therefore suppresses power on scales larger than the half-mode. Instead, our criterion for imposing the mass bound on quasithermal PSDs follows the criterion established in Ref. [99]. This criterion is further supported by Ref. [34], which shows near-perfect agreement between their area-based criterion, a more sophisticated integrated power criterion, and the warmness bound derived from the second moment σ_q of the DM PSD. The origin of the discrepancy between these two mass bounds is thus explained: achieving an allowed power spectrum requires a larger high- k cutoff, which in turn implies a higher $m_\chi^{\min}$.³

VIII. CONCLUSIONS

The freeze-in mechanism is an appealing framework with many merits, though it remains highly elusive to experimental searches. In this work, we investigated its potential imprints on cosmological observables at small scales. The reason why DM produced via freeze-in can leave such signatures lies in a key feature of light DM with (quasi)thermal origin: its free-streaming velocities can be

large enough to suppress the formation of the smallest structures. Since this suppression depends sensitively on the PSD of freeze-in DM, accurately predicting this distribution through the Boltzmann equation is essential.

A simple and efficient way to assess whether small-scale structures are preserved is to compare freeze-in DM with a quasithermal PSD to the well-studied case of thermal WDM, which serves as a benchmark for the impact of nonzero free-streaming velocities. In this work, we developed a general methodology based on the criterion introduced in Ref. [31] and numerically validated in Ref. [34], where the key mapping variable is the root-mean-square velocity, $W_\chi = \sqrt{\langle p^2 \rangle / m_\chi}$. This approach allowed us to bypass the need for computationally intensive cosmological Boltzmann solvers such as CLASS [87] when scanning the parameter space of explicit freeze-in models. Thanks to this strategy, we were able to extract lower bounds on the DM mass across a wide range of scenarios.

The plot in Fig. 7 offers a compact summary of all the results presented in this paper. The two axes show the dimensionless second moment of the PSD, σ_q [defined in Eq. (6)], and the temperature scale T_p introduced in Eq. (2) to define the comoving momentum. Solid black lines indicate the lower bounds on the DM mass obtained via the rescaling in Eq. (7). For this figure, we adopt the most stringent constraint from MW satellite counts, $m_{\text{WDM}}^{\min} = 6.8$ keV. As explained in the main text, these bounds are independent of the *arbitrary* choice of T_p , which can be selected conveniently. It is natural to set T_p to the largest mass involved in the production process, a choice we adopt for all microscopic realizations analyzed in this work. In the case of the dark photon portal, where both plasmon decay and electron annihilation contribute, T_p is defined as the temperature at which the plasmon frequency equals $2m_\gamma$, which exceeds the electron mass. Once T_p is fixed, σ_q is correspondingly determined, and each model appears as a single point in the (σ_q, T_p) plane. All models discussed in the paper are represented this way, as indicated in the legend. The minimum allowed DM mass can then be directly inferred from the location of the point relative to the black isocontours.

The usefulness of presenting the results in this format lies in the fact that it enables a rapid visualization of different models and provides insight into which freeze-in DM candidates can be considered “colder” and which “warmer.” For instance, ALPs produced via scatterings with electrons and DM produced through the dark photon portal were both assigned similar values of T_p , as production in these scenarios naturally occurs at comparable energy scales. However, the corresponding σ_q differ significantly, resulting in DM from the dark photon portal being much colder than that from ALP-electron scatterings. This distinction arises from the different production mechanisms and associated kinematics in the two models. The roles of T_p and σ_q are clearly illustrated: T_p serves as a

³We have explicitly verified that our bound is consistent with results obtained by inputting the PSD into the Boltzmann solver CLASS [87] and computing the power suppression at small scales.




FIG. 7. This figure summarizes all the results obtained in this paper. The variable on the horizontal axis, σ_q , is defined in Eq. (6) as the second moment of the PSD evaluated in terms of the comoving momentum. It quantifies the width of the momentum-space distribution. The vertical axis shows the variable T_P , which enters the definition of the comoving momentum in Eq. (2). While T_P is a conventional parameter, it is convenient to identify it with the temperature at which freeze-in is most efficient—typically corresponding to the mass of the heaviest bath particle participating in DM production. The black isocontours (identical to those in Fig. 3) represent the minimum allowed DM mass produced via freeze-in. These values are obtained through the rescaling given in Eq. (7), using $m_{\text{WDM}}^{\text{min}} = 6.8$ keV as a reference. Each model realization is indicated in the plot by the symbols shown in the legend. In each case, T_P is chosen appropriately as the mass of the heaviest particle involved in the decay or scattering process that produces DM. For the dark photon portal, T_P is defined by the condition $\omega_p(T_P) = 2m_\chi$. For the ALP DM scatterings we show the bounds obtained by assuming constant matrix elements as empty circles. We show the sterile neutrino model as a dashed cyan line, since the mass of the decaying particle is not specified, therefore T_P is also varying; the Higgs portal case is shown by a pink cross, the dark photon portal as a gold star. The dots refer to ALP DM produced by flavor conserving (FC) scatterings of SM fermions with gauge bosons, while the pluses to ALP DM produced by flavor violating (FV) decays of SM fermions. Blue markers indicate processes with the tau lepton, orange with the muon, green with the electron, red with the top quark, violet with the bottom, and brown with the charm.

convenient reference scale, typically set by the dominant energy scale of the production process, while σ_q , once defined with respect to this scale, quantifies the “warmness” of the resulting DM population.

Another important result of this work becomes manifest when comparing the model-independent analysis presented in Fig. 4 with the summary plot shown in Fig. 7. The former was obtained under the assumptions of constant matrix elements for all production processes and Maxwell-Boltzmann statistics for all bath particles. In contrast, for each microscopic realization, we evaluated the transition amplitudes self-consistently and included quantum statistical effects appropriate to the particles involved. For the ALP DM produced via binary scatterings, we show the bounds on the DM mass obtained under the constant matrix element—that has been employed in setting model-independent bounds in Sec. IV—as empty circles. Remarkably, the resulting model points align with the model-independent predictions up to 5%–10% accuracy,

with the constant matrix element assumption providing a more conservative bound. An important exception is the dark photon portal, for which the assumption of a single production mechanism does not hold due to contamination from plasmon decays. This agreement is not unexpected: IR ensures that the momentum dependence of the matrix elements has only a minor impact, and Ref. [34] explicitly checked that quantum degeneracy effects can be safely neglected for freeze-in production in the early Universe. These findings validate the robustness of our model-independent mass bounds under the assumption of no interferences between production mechanisms.

To summarize, this work provided an explicit and computationally efficient methodology to derive mass bounds on freeze-in DM in a model-independent manner. The results presented here lay the foundation for future studies involving DM particles produced with a quasithermal momentum distribution. Notable examples include scenarios with a subdominant DM component [34,121],

UV-dominated freeze-in production [26,33,122–131], or freeze-in from a decoupled thermal bath with a temperature different from that of the SM degrees of freedom [132]. We leave these directions for future investigation.

ACKNOWLEDGMENTS

We thank Mathias Becker for useful discussions. F. D. is supported by Istituto Nazionale di Fisica Nucleare (INFN) through the Theoretical Astroparticle Physics (TAsP) project. This work is supported in part by the Italian MUR Departments of Excellence Grant No. 2023-2027 “Quantum Frontiers.” The work of A. L. is supported by an ERC STG grant (“Light-Dark,” Grant No. 101040019). This project has received funding from the European Research Council (ERC) under the European Union’s Horizon Europe research and innovation programme (Grant Agreement No. 101040019). A. L. is grateful to the Azrieli Foundation for the award of an Azrieli Fellowship and to the Department of Physics and Astronomy in Padua for the hospitality during the completion of this work. This article is based upon work from COST Action COSMIC WISPerS CA21106, supported by COST (European Cooperation in Science and Technology). A. D. is supported by the Kavli Institute for Cosmological physics at the University of Chicago through an endowment from the Kavli Foundation and its founder Fred Kavli.

Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union. The European Union cannot be held responsible for them.

DATA AVAILABILITY

The data are not publicly available. The data are available from the authors upon reasonable request.

APPENDIX A: THE WDM PARADIGM

In this appendix, we review the WDM framework. The central assumption is that DM is distributed in momentum space according to a thermal PSD. We define T_{WDM} as the current temperature of the WDM population, which may differ from the present CMB temperature T_0 . The PSD takes the explicit form

$$f_{\text{WDM}}(p) = \begin{cases} [\exp(p/T_{\text{WDM}}) - 1]^{-1} & \text{(BE)} \\ [\exp(p/T_{\text{WDM}}) + 1]^{-1} & \text{(FD)} \\ \exp(-p/T_{\text{WDM}}) & \text{(MB)} \end{cases}, \quad (\text{A1})$$

corresponding to Bose-Einstein (BE), Fermi-Dirac (FD), and Maxwell-Boltzmann (MB) statistics. The framework is specified by two parameters: the DM mass m_{WDM} and its

temperature T_{WDM} , which determines the shape of the distribution independently of the thermal bath.

For a WDM candidate with g_{WDM} internal degrees of freedom (e.g., $g_{\text{WDM}} = 2$ for a Majorana fermion), the DM number density is obtained by integrating the phase-space distribution

$$\begin{aligned} n_{\text{WDM}}(t_0) &= g_{\text{WDM}} \int \frac{d^3 p}{(2\pi)^3} f_{\text{WDM}}(p) \\ &= g_{\text{WDM}} g_{\text{eff}} \frac{\zeta(3)}{\pi^2} T_{\text{WDM}}^3, \end{aligned} \quad (\text{A2})$$

where $\zeta(3)$ is the Riemann zeta function. The statistical factor $g_{\text{eff}} = \{1, 3/4, 1/\zeta(3)\}$ arises from the momentum integral of the PSD in the BE, FD, and MB cases, respectively, as listed in Eq. (A1).

The current WDM energy density is given by multiplying the number density at the present time t_0 by the WDM mass, $\rho_{\text{WDM}}(t_0) = m_{\text{WDM}} n_{\text{WDM}}(t_0)$. It is convenient to express this quantity in terms of the usual parameter Ω_{WDM} , defined as

$$\Omega_{\text{WDM}} h^2 \equiv \frac{\rho_{\text{WDM}}(t_0)}{\rho_{\text{cr}}(t_0)/h^2}, \quad (\text{A3})$$

where $\rho_{\text{cr}}(t_0) \simeq 1.05 \times 10^{-5} h^2 \text{ GeV cm}^{-3}$ is the current critical density [133]. Combining the above expressions, the WDM temperature can be related to the present CMB temperature T_0 as follows

$$\frac{T_{\text{WDM}}}{T_0} \simeq 0.16 \left(\frac{2}{g_{\text{WDM}}} \frac{3/4}{g_{\text{eff}}} \frac{\text{keV}}{m_{\text{WDM}}} \frac{\Omega_{\text{WDM}} h^2}{0.12} \right)^{1/3}. \quad (\text{A4})$$

Here, we have normalized the expression using parameters appropriate for a Majorana fermion, as WDM mass bounds are typically quoted under this assumption.

At this point, we can evaluate the quantity defined in Eq. (5), which characterizes the WDM warmth

$$W_{\text{WDM}} = \frac{\sqrt{\langle p^2 \rangle}}{m_{\text{WDM}}} = \sigma_q^{\text{WDM}} \frac{T_{\text{WDM}}}{m_{\text{WDM}}}, \quad (\text{A5})$$

where σ_q^{WDM} is the dimensionless second moment of the WDM distribution, given by

$$\sigma_q^{\text{WDM}} = \begin{cases} 2\sqrt{3\zeta(5)/\zeta(3)} \simeq 3.22 & \text{(BE)} \\ \sqrt{15\zeta(5)/\zeta(3)} \simeq 3.60 & \text{(FD)} \\ 2\sqrt{3} \simeq 3.46 & \text{(MB)} \end{cases}. \quad (\text{A6})$$

APPENDIX B: COLLISION OPERATORS AND ANALYTIC PSDs

We summarize the collision operators for the topologies used in this work and the analytical expressions for the PSD. These and other related results can be found in Ref. [34] with the exception of the new analytical PSD for three-body decays given in Eq. (B16). We work in the light DM limit, $m_\chi \ll m_1$, and we report the expressions for DM single production. As discussed in the text, these results are valid also for multiple production unless we have almost degenerate bath particles. The analytical expressions of the PSD provided here match the numerical results to high accuracy in the limit $g_* = \text{const}$.

1. Two-body decays

For a two-body decay the collision term reads

$$C_{1 \rightarrow 2\chi} = \frac{\ell}{2} \int d\Pi_1 d\Pi_2 (2\pi)^4 \delta^{(4)}(P_1 - P_2 - P) |\mathcal{M}|^2 f_1^{\text{eq}}. \quad (\text{B1})$$

The matrix element is always constant, and the phase space integration leads to the following result

$$C_{1 \rightarrow 2\chi}(p, T) = \frac{\ell g_1 g_2 T}{16\pi p} |\mathcal{M}|^2 h(p, T). \quad (\text{B2})$$

The dimensionless function $h(p, T)$ depends on the statistics of the decaying bath particle

$$h(p, T) \approx \begin{cases} e^{-(E_2^- + E)/T} - e^{-(E_2^+ + E)/T} & (\text{MB}) \\ -\log \left[\frac{1 - e^{-(E_2^- + E)/T}}{1 - e^{-(E_2^+ + E)/T}} \right] & (\text{BE}) \\ \log \left[\frac{1 + e^{-(E_2^- + E)/T}}{1 + e^{-(E_2^+ + E)/T}} \right] & (\text{FD}) \end{cases} \quad (\text{B3})$$

and the bounds on E_2 being $E_2^+ \rightarrow \infty$

$$E_2^- \approx \frac{m_1^2}{4pT} \left(1 - \frac{m_2^2}{m_1^2} \right) + \frac{pm_2^2}{m_1^2 - m_2^2} + \mathcal{O}(m_\chi). \quad (\text{B4})$$

An analytical PSD estimate can be obtained by integrating the collision term assuming $g_* = \text{const}$, radiation domination, MB statistics, and $q = p/T$. We find

$$f_\chi(q) \propto \frac{1}{\sqrt{q}} e^{-q(1 - m_2^2/m_1^2)^{-1}}. \quad (\text{B5})$$

This implies $\sigma_q \simeq 2.958 \sqrt{1 - (m_2/m_1)^2}$. Note how the phase space suppression m_2/m_1 makes the PSD “colder,” reducing the momentum dispersion.

2. Binary scatterings

For a binary scattering we have

$$C_{12 \rightarrow 3\chi} = \frac{\ell}{2} \int \prod_{i=1}^3 d\Pi_i (2\pi)^4 \delta^{(4)}(P_i - P_f) |\mathcal{M}|^2 f_1^{\text{eq}} f_2^{\text{eq}}. \quad (\text{B6})$$

The expression greatly simplifies for MB statistics

$$C_{12 \rightarrow 3\chi} = \frac{\ell g_1 g_2 g_3 T e^{-E/T}}{256\pi^3 p} \times \int_{s_-}^{\infty} \frac{ds e^{-E_3^-/T}}{\sqrt{\lambda(s, m_3, m_\chi)}} \int_{t_-}^{t_+} dt |\mathcal{M}|^2, \quad (\text{B7})$$

we have introduced the Källén function $\lambda(a, b, c) \equiv [a - (b + c)^2][a - (b - c)^2]$. Here $E_3^- \approx \sqrt{m_3^2 + (s/4p)^2} + \mathcal{O}(m_\chi)$ and the integration limits read

$$s_- = \max[(m_1 + m_2)^2, (m_3 + m_\chi)^2] \quad (\text{B8})$$

$$t_\pm = m_2^2 + m_\chi^2 - 2 \left(\sqrt{m_2^2 + p_{12}^2} \sqrt{m_3^2 + p_{3\chi}^2} \mp p_{12} p_{3\chi} \right) \quad (\text{B9})$$

where $p_{ij}(s) = \sqrt{\lambda(s, m_i, m_j)}/2$.

The matrix element is not constant in general for this case. If we approximate it to a constant value, and with the further assumptions of constant g_* , radiation domination, MB statistics, $q = p/T$, and $m_2 = m_3 = 0$, we find the analytical solution

$$f_\chi(q) \propto \int \frac{dT e^{-p/T}}{T^2 H p^2} \int_{m_1^2}^{\infty} ds \frac{s-1}{s} e^{-s/4pT} \propto \frac{1}{\sqrt{q}} e^{-q}. \quad (\text{B10})$$

This implies $\sigma_q \simeq 2.958$.

3. Three-body decays

Finally, for three-body decays we have

$$C_{1 \rightarrow 23\chi} = \frac{\ell}{2} \int \prod_{i=1}^3 d\Pi_i (2\pi)^4 \delta^{(4)}(P_i - P_f) |\mathcal{M}|^2 f_1^{\text{eq}}. \quad (\text{B11})$$

Also in this case the matrix element is not constant and has to be integrated over specially defined Mandelstam-like variables $s = (P_1 - P)^2$ and $t = (P_1 - P_2)^2$. If we assume MB statistics

$$C_{1 \rightarrow 23\chi} = \frac{\ell g_1 g_2 g_3 T}{256\pi^3 p} \int_{s_-}^{s_+} \frac{ds e^{-E_1^-/T}}{\sqrt{\lambda(s, m_1, m_\chi)}} \int_{t_-}^{t_+} dt |\mathcal{M}|^2. \quad (\text{B12})$$

The integration limits read

$$s_+ = (m_1 - m_\chi)^2, \quad s_- = (m_2 + m_3)^2, \quad (\text{B13})$$

$$t_\pm = m_3^2 + m_\chi^2 + 2\left(\sqrt{m_3^2 + p_{23}^2}\sqrt{m_\chi^2 + p_{1\chi}^2} \pm p_{1\chi}p_{23}\right), \quad (\text{B14})$$

where $p_{ij}(s) = \sqrt{\lambda(s, m_i, m_j)}/s/2$. E_1^- is given as

$$E_1^- \approx \frac{p}{1 - s/m_1^2} + \frac{m_1^2 - s}{4p} + \mathcal{O}(m_\chi). \quad (\text{B15})$$

We can obtain (a rather peculiar) analytical form for the PSD. The assumptions are again constant g_* , radiation domination, MB statistics, $|\mathcal{M}|^2$ constant, $q = p/T$, $m_3 = 0$. We keep m_2 nonzero to show the phase space

suppression, and we find

$$f_\chi(q) \propto \int \frac{dT}{H} \frac{1}{p^2} \int_{m_2^2}^{m_1^2} ds e^{-E_1^-/T} \left(1 - \frac{m_2^2}{s}\right) \propto \frac{1}{q} \text{erfc}\left(\sqrt{\frac{q}{1 - m_2^2/m_1^2}}\right) \quad (\text{B16})$$

In this case the integral over temperatures was performed before the one over s . Remarkably, we obtained an expression that takes into account the phase space suppression. We obtain $\sigma_q \simeq 2.092\sqrt{1 - (m_2/m_1)^2}$, showing how the mass of particle 2 makes the PSD colder. Note also how, even for $m_2 = 0$, three-body decays have a smaller σ_q than two-body decays.

-
- [1] G. Jungman, M. Kamionkowski, and K. Griest, *Phys. Rep.* **267**, 195 (1996).
 - [2] G. Bertone, D. Hooper, and J. Silk, *Phys. Rep.* **405**, 279 (2005).
 - [3] J. L. Feng, *Annu. Rev. Astron. Astrophys.* **48**, 495 (2010).
 - [4] P. Salucci, *Astron. Astrophys. Rev.* **27**, 2 (2019).
 - [5] M. Cirelli, A. Strumia, and J. Zupan, arXiv:2406.01705.
 - [6] J. McDonald, *Phys. Rev. Lett.* **88**, 091304 (2002).
 - [7] A. Kusenko, *Phys. Rev. Lett.* **97**, 241301 (2006).
 - [8] A. Ibarra, A. Ringwald, and C. Weniger, *J. Cosmol. Astropart. Phys.* **01** (2008) 003.
 - [9] L. J. Hall, K. Jedamzik, J. March-Russell, and S. M. West, *J. High Energy Phys.* **03** (2009) 080.
 - [10] R. T. Co, F. D'Eramo, L. J. Hall, and D. Pappadopulo, *J. Cosmol. Astropart. Phys.* **12** (2015) 024.
 - [11] J. A. Evans and J. Shelton, *J. High Energy Phys.* **04** (2016) 056.
 - [12] F. D'Eramo, N. Fernandez, and S. Profumo, *J. Cosmol. Astropart. Phys.* **02** (2017) 046.
 - [13] L. Calibbi, L. Lopez-Honorez, S. Lowette, and A. Mariotti, *J. High Energy Phys.* **09** (2018) 037.
 - [14] D. Curtin *et al.*, *Rep. Prog. Phys.* **82**, 116201 (2019).
 - [15] G. Bélanger *et al.*, *J. High Energy Phys.* **02** (2018) 186.
 - [16] S. Junius, L. Lopez-Honorez, and A. Mariotti, *J. High Energy Phys.* **07** (2019) 136.
 - [17] J. M. No, P. Tunney, and B. Zaldivar, *J. High Energy Phys.* **03** (2019) 022.
 - [18] K. J. Bae, M. Park, and M. Zhang, *Phys. Rev. D* **101**, 115036 (2020).
 - [19] L. Calibbi, F. D'Eramo, S. Junius, L. Lopez-Honorez, and A. Mariotti, *J. High Energy Phys.* **05** (2021) 234.
 - [20] X. Chu, T. Hambye, and M. H. G. Tytgat, *J. Cosmol. Astropart. Phys.* **05** (2011) 034.
 - [21] R. Essig, M. Fernandez-Serra, J. Mardon, A. Soto, T. Volansky, and T.-T. Yu, *J. High Energy Phys.* **05** (2015) 046.
 - [22] C. Dvorkin, T. Lin, and K. Schutz, *Phys. Rev. D* **99**, 115009 (2019); **105**, 119901(E) (2022).
 - [23] K. K. Boddy, K. Freese, G. Montefalcone, and B. Shams Es Haghi, *Phys. Rev. D* **111**, 063537 (2025).
 - [24] N. Bernal, C. S. Fong, and Ó. Zapata, *J. High Energy Phys.* **02** (2024) 161.
 - [25] A. Kamada, N. Yoshida, K. Kohri, and T. Takahashi, *J. Cosmol. Astropart. Phys.* **03** (2013) 008.
 - [26] J. McDonald, *J. Cosmol. Astropart. Phys.* **08** (2015) 035.
 - [27] S. B. Roland and B. Shakya, *J. Cosmol. Astropart. Phys.* **05** (2016) 027.
 - [28] J. Heeck and D. Teresi, *Phys. Rev. D* **96**, 035018 (2017).
 - [29] K. J. Bae, A. Kamada, S. P. Liew, and K. Yanagi, *J. Cosmol. Astropart. Phys.* **01** (2017) 054.
 - [30] S. Boulebnane, J. Heeck, A. Nguyen, and D. Teresi, *J. Cosmol. Astropart. Phys.* **04** (2017) 006.
 - [31] A. Kamada and K. Yanagi, *J. Cosmol. Astropart. Phys.* **11** (2019) 029.
 - [32] C. Dvorkin, T. Lin, and K. Schutz, *Phys. Rev. Lett.* **127**, 111301 (2021).
 - [33] G. Ballesteros, M. A. G. Garcia, and M. Pierre, *J. Cosmol. Astropart. Phys.* **03** (2020) 101.
 - [34] F. D'Eramo and A. Lenoci, *J. Cosmol. Astropart. Phys.* **10** (2020) 045.
 - [35] S. Baumholzer, V. Brdar, and E. Morgante, *J. Cosmol. Astropart. Phys.* **05** (2020) 004.
 - [36] D. Egana-Ugrinovic, R. Essig, D. Gift, and M. LoVerde, *J. Cosmol. Astropart. Phys.* **05** (2021) 013.
 - [37] Y. Du, F. Huang, H.-L. Li, Y.-Z. Li, and J.-H. Yu, *J. Cosmol. Astropart. Phys.* **04** (2021) 012.
 - [38] Q. Decant, J. Heisig, D. C. Hooper, and L. Lopez-Honorez, *J. Cosmol. Astropart. Phys.* **03** (2021) 041.
 - [39] K. R. Dienes, F. Huang, J. Kost, B. Thomas, and H.-B. Yu, *Phys. Rev. D* **106**, 123521 (2022).
 - [40] Z. Xu, Q. Zhou, and S. Zheng, *Phys. Rev. D* **110**, 115003 (2024).

- [41] J. S. Bullock and M. Boylan-Kolchin, *Annu. Rev. Astron. Astrophys.* **55**, 343 (2017).
- [42] K. Bechtol *et al.*, in Snowmass 2021 (2022), [arXiv:2203.07354](#).
- [43] A. Dekker, S. Ando, C. A. Correa, and K. C. Y. Ng, *Phys. Rev. D* **106**, 123026 (2022).
- [44] E. O. Nadler *et al.* (DES Collaboration), *Phys. Rev. Lett.* **126**, 091101 (2021).
- [45] O. Newton, M. Leo, M. Cautun, A. Jenkins, C. S. Frenk, M. R. Lovell, J. C. Helly, A. J. Benson, and S. Cole, *J. Cosmol. Astropart. Phys.* **08** (2020) 062.
- [46] R. Kennedy, C. Frenk, S. Cole, and A. Benson, *Mon. Not. R. Astron. Soc.* **442**, 2487 (2014).
- [47] M. R. Lovell, C. S. Frenk, V. R. Eke, A. Jenkins, L. Gao, and T. Theuns, *Mon. Not. R. Astron. Soc.* **439**, 300 (2014).
- [48] A. Schneider, *Mon. Not. R. Astron. Soc.* **451**, 3117 (2015).
- [49] E. Polisensky and M. Ricotti, *Phys. Rev. D* **83**, 043506 (2011).
- [50] M. Cautun, A. Benítez-Llambay, A. J. Deason, C. S. Frenk, A. Fattahi, F. A. Gómez, R. J. J. Grand, K. A. Oman, J. F. Navarro, and C. M. Simpson, *Mon. Not. R. Astron. Soc.* **494**, 4291 (2020).
- [51] T. Callingham, M. Cautun, A. J. Deason, C. S. Frenk, W. Wang, F. A. Gómez, R. J. J. Grand, F. Marinacci, and R. Pakmor, *Mon. Not. R. Astron. Soc.* **484**, 5453 (2019).
- [52] G. Eadie and M. Jurić, *Astrophys. J.* **875**, 159 (2019).
- [53] L. Posti and A. Helmi, *Astron. Astrophys.* **621**, A56 (2019).
- [54] E. Karukes, M. Benito, F. Iocco, R. Trotta, and A. Geringer-Sameth, *J. Cosmol. Astropart. Phys.* **05** (2020) 033.
- [55] J. Wang, F. Hammer, and Y. Yang, *Mon. Not. R. Astron. Soc.* **510**, 2242 (2021).
- [56] S. A. Bird, X.-X. Xue, C. Liu, C. Flynn, J. Shen, J. Wang, C. Yang, M. Zhai, L. Zhu, G. Zhao, and H.-J. Tian, *Mon. Not. R. Astron. Soc.* **516**, 731 (2022).
- [57] E. O. Nadler *et al.* (DES Collaboration), *Astrophys. J.* **893**, 48 (2020).
- [58] V. V. Bobylev and A. T. Baykova, *Astronomy Reports* **67**, 812 (2023).
- [59] A. M. Nierenberg *et al.*, *Mon. Not. R. Astron. Soc.* **530**, 2960 (2024).
- [60] S. Vegetti *et al.*, *Space Sci. Rev.* **220**, 58 (2024).
- [61] J.-W. Hsueh, W. Enzi, S. Vegetti, M. Auger, C. D. Fassnacht, G. Despali, L. V. E. Koopmans, and J. P. McKean, *Mon. Not. R. Astron. Soc.* **492**, 3047 (2020).
- [62] D. Gilman, S. Birrer, A. Nierenberg, and M. S. H. Oh, *Mon. Not. R. Astron. Soc.* **533**, 1687 (2024).
- [63] R. E. Keeley *et al.*, *Mon. Not. R. Astron. Soc.* **535**, 1652 (2024).
- [64] A. Garzilli, A. Magalich, T. Theuns, C. S. Frenk, C. Weniger, O. Ruchayskiy, and A. Boyarsky, *Mon. Not. R. Astron. Soc.* **489**, 3456 (2019).
- [65] A. Garzilli, A. Magalich, O. Ruchayskiy, and A. Boyarsky, *Mon. Not. R. Astron. Soc.* **502**, 2356 (2021).
- [66] J. Baur, N. Palanque-Delabrouille, C. Yèche, C. Magneville, and M. Viel, *J. Cosmol. Astropart. Phys.* **08** (2015) 012.
- [67] A. Garzilli, A. Boyarsky, and O. Ruchayskiy, *Phys. Lett. B* **773**, 258 (2017).
- [68] B. Villasenor, B. Robertson, P. Madau, and E. Schneider, *Phys. Rev. D* **108**, 023502 (2023).
- [69] V. Iršič *et al.*, *Phys. Rev. D* **109**, 043511 (2024).
- [70] J. Bovy, D. Erkal, and J. L. Sanders, *Mon. Not. R. Astron. Soc.* **466**, 628 (2017).
- [71] D. Erkal and V. Belokurov, *Mon. Not. R. Astron. Soc.* **454**, 3542 (2015).
- [72] N. Banik, H.-J. He, and E. D. Schiappacasse, *J. Cosmol. Astropart. Phys.* **10** (2019) 043.
- [73] N. Banik, J. Bovy, G. Bertone, D. Erkal, and T. J. L. de Boer, *Mon. Not. R. Astron. Soc.* **502**, 2364 (2021).
- [74] R. Ibata, G. Thomas, B. Famaey, K. Malhan, N. Martin, and G. Monari, *Astrophys. J.* **891**, 161 (2020).
- [75] A. Bonaca and A. M. Price-Whelan, *New Astron. Rev.* **100**, 101713 (2025).
- [76] T. Hilmi *et al.* (S5 Collaboration), [arXiv:2404.02953](#).
- [77] R. G. Carlberg, A. Jenkins, C. S. Frenk, and A. P. Cooper, *Astrophys. J.* **975**, 135 (2024).
- [78] A. Drlica-Wagner *et al.* (LSST Dark Matter Group Collaboration), [arXiv:1902.01055](#).
- [79] P. Corasaniti, S. Agarwal, D. Marsh, and S. Das, *Phys. Rev. D* **95**, 083512 (2017).
- [80] C. Schultz, J. Oñorbe, K. N. Abazajian, and J. S. Bullock, *Mon. Not. R. Astron. Soc.* **442**, 1597 (2014).
- [81] N. Menci, A. Grazian, M. Castellano, and N. G. Sanchez, *Astrophys. J. Lett.* **825**, L1 (2016).
- [82] J. Ellis, M. Fairbairn, J. Urrutia, and V. Vaskonen, *Astron. Astrophys.* **702**, A109 (2025).
- [83] B. Liu, H. Shan, and J. Zhang, *Astrophys. J.* **968**, 79 (2024).
- [84] U. Maio and M. Viel, *Astron. Astrophys.* **672**, A71 (2023).
- [85] W. Enzi *et al.*, *Mon. Not. R. Astron. Soc.* **506**, 5848 (2021).
- [86] E. O. Nadler, S. Birrer, D. Gilman, R. H. Wechsler, X. Du, A. Benson, A. M. Nierenberg, and T. Treu, *Astrophys. J.* **917**, 7 (2021).
- [87] J. Lesgourgues and T. Tram, *J. Cosmol. Astropart. Phys.* **09** (2011) 032.
- [88] M. Laine and M. Meyer, *J. Cosmol. Astropart. Phys.* **07** (2015) 035.
- [89] C. P. Burgess, M. Pospelov, and T. ter Veldhuis, *Nucl. Phys.* **B619**, 709 (2001).
- [90] J. McDonald, *Phys. Rev. D* **50**, 3637 (1994).
- [91] N. Bernal, C. Cosme, T. Tenkanen, and V. Vaskonen, *Eur. Phys. J. C* **79**, 30 (2019).
- [92] F. D'ErAMO, A. Tesi, and V. Vaskonen, *Phys. Rev. D* **110**, 095002 (2024).
- [93] S. Dodelson and L. M. Widrow, *Phys. Rev. Lett.* **72**, 17 (1994).
- [94] C. Dessert, J. W. Foster, Y. Park, and B. R. Safdi, *Astrophys. J.* **964**, 185 (2024).
- [95] K. N. Abazajian, *Phys. Rep.* **711–712**, 1 (2017).
- [96] X.-D. Shi and G. M. Fuller, *Phys. Rev. Lett.* **82**, 2832 (1999).
- [97] A. De Gouvêa, M. Sen, W. Tangarife, and Y. Zhang, *Phys. Rev. Lett.* **124**, 081802 (2020).
- [98] K. J. Kelly, M. Sen, W. Tangarife, and Y. Zhang, *Phys. Rev. D* **101**, 115031 (2020).
- [99] J. König, A. Merle, and M. Totzauer, *J. Cosmol. Astropart. Phys.* **11** (2016) 038.

- [100] R. D. Peccei and H. R. Quinn, *Nuovo Cimento A* **41**, 309 (1977).
- [101] R. D. Peccei and H. R. Quinn, *Phys. Rev. D* **16**, 1791 (1977).
- [102] F. Wilczek, *Phys. Rev. Lett.* **40**, 279 (1978).
- [103] S. Weinberg, *Phys. Rev. Lett.* **40**, 223 (1978).
- [104] F. D’Eramo, R. Z. Ferreira, A. Notari, and J. L. Bernal, *J. Cosmol. Astropart. Phys.* **11** (2018) 014.
- [105] M. Badziak and M. Laletin, *J. High Energy Phys.* **02** (2024) 108.
- [106] F. D’Eramo and A. Lenoci, *Phys. Rev. D* **110**, 116028 (2024).
- [107] F. Arias-Aragón, F. D’Eramo, R. Z. Ferreira, L. Merlo, and A. Notari, *J. Cosmol. Astropart. Phys.* **03** (2020) 090.
- [108] F. D’Eramo, F. Hajkarim, and A. Lenoci, *J. Cosmol. Astropart. Phys.* **03** (2023) 009.
- [109] M. Bauer, M. Neubert, and A. Thamm, *J. High Energy Phys.* **12** (2017) 044.
- [110] M. Bauer, M. Neubert, S. Renner, M. Schnubel, and A. Thamm, *J. High Energy Phys.* **04** (2020) 063.
- [111] D. Cadamuro and J. Redondo, *J. Cosmol. Astropart. Phys.* **02** (2011) 032.
- [112] K. Perez, K. C. Y. Ng, J. F. Beacom, C. Hersh, S. Horiuchi, and R. Krivonos, *Phys. Rev. D* **95**, 123002 (2017).
- [113] K. C. Y. Ng, B. M. Roach, K. Perez, J. F. Beacom, S. Horiuchi, R. Krivonos, and D. R. Wik, *Phys. Rev. D* **99**, 083005 (2019).
- [114] B. M. Roach, S. Rossland, K. C. Y. Ng, K. Perez, J. F. Beacom, B. W. Grefenstette, S. Horiuchi, R. Krivonos, and D. R. Wik, *Phys. Rev. D* **107**, 023009 (2023).
- [115] P. Panci, D. Redigolo, T. Schwetz, and R. Ziegler, *Phys. Lett. B* **841**, 137919 (2023).
- [116] M. Aghaie, G. Armando, A. Conaci, A. Dondarini, P. Maták, P. Panci, Z. Šinská, and R. Ziegler, *Phys. Lett. B* **856**, 138923 (2024).
- [117] F. D’Eramo and S. Yun, *Phys. Rev. D* **105**, 075002 (2022).
- [118] J. H. Chang, R. Essig, and A. Reinert, *J. High Energy Phys.* **03** (2019) 141.
- [119] P. N. Bhattiprolu, R. McGehee, E. Petrosky, and A. Pierce, *Phys. Rev. D* **111**, 035027 (2025).
- [120] R. Murgia, V. Iršič, and M. Viel, *Phys. Rev. D* **98**, 083540 (2018).
- [121] E. Iles, S. Heeba, and K. Schutz, *Phys. Rev. Lett.* **134**, 121002 (2025).
- [122] F. Elahi, C. Kolda, and J. Unwin, *J. High Energy Phys.* **03** (2014) 048.
- [123] S.-L. Chen and Z. Kang, *J. Cosmol. Astropart. Phys.* **05** (2017) 036.
- [124] N. Bernal, F. Elahi, C. Maldonado, and J. Unwin, *J. Cosmol. Astropart. Phys.* **11** (2019) 026.
- [125] B. Barman, N. Bernal, Y. Xu, and Ó. Zapata, *J. Cosmol. Astropart. Phys.* **07** (2022) 019.
- [126] A. Ahmed, B. Grzadkowski, and A. Socha, *J. High Energy Phys.* **02** (2022) 196.
- [127] D. K. Ghosh, A. Ghoshal, and S. Jeesun, *J. High Energy Phys.* **01** (2023) 026.
- [128] K. Freese, G. Montefalcone, and B. Shams Es Haghi, *Phys. Rev. Lett.* **133**, 211001 (2024).
- [129] R. de Souza, J. G. Rodrigues, C. Siqueira, F. B. M. d. Santos, and J. Alcaniz, *J. High Energy Phys.* **04** (2024) 125.
- [130] L. Caloni, P. Stengel, M. Lattanzi, and M. Gerbino, *J. Cosmol. Astropart. Phys.* **10** (2024) 106.
- [131] N. Bernal, K. Deka, and M. Losada, *Phys. Rev. D* **111**, 055034 (2025).
- [132] N. Fernandez, Y. Kahn, and J. Shelton, *J. High Energy Phys.* **07** (2021) 044.
- [133] S. Navas *et al.* (Particle Data Group), *Phys. Rev. D* **110**, 030001 (2024).