Electroweak axion portal to dark matter

Stephanie Allen,[1] Albany Blackburn,[1,2] Oswaldo Cardenas,[1] Zoe Messenger,[1] Ngan H. Nguyen,[1,3] and Brian Shuve[1]

[1]*Harvey Mudd College, 301 Platt Blvd., Claremont, California 91711, USA*
[2]*Berkeley Center for Theoretical Physics, University of California, Berkeley, California 94720, USA*
[3]*Johns Hopkins University, 3400 N Charles St., Baltimore, Maryland 21218, USA*

(Received 3 June 2024; accepted 6 October 2024; published 12 November 2024)

Axion-like particles (ALPs) are good candidates for mediators to the dark sector. We explore scenarios in which an ALP mediates interactions between dark matter and electroweak gauge bosons. These models yield testable electromagnetic signals in astrophysical, cosmological, and terrestrial probes. We find promising prospects for both indirect detection and accelerator tests, with interesting parameter space already constrained by current experiments. Our work provides concrete benchmarks for future tests of the electroweak ALP portal.

DOI: 10.1103/PhysRevD.110.095010

I. INTRODUCTION

Axion-like particles (ALPs) are among the best-motivated candidates for new particles beyond the Standard Model (SM). They arise as pseudo-Goldstone bosons whenever there exist spontaneously broken (approximate) global symmetries that are anomalous under SM gauge interactions, also known as Peccei-Quinn (PQ) symmetries [1–4]. ALPs are generic features of ultraviolet (UV) theories such as string theory [5–8] and supersymmetry [9–11]. Although axions were originally hypothesized to explain the absence of CP-violation in quantum chromodynamics (QCD) [1–4], ALPs have also been proposed to play a role in explaining the smallness of the electroweak scale relative to the Planck scale [12].

While the PQ symmetry is spontaneously broken at some energy scale f, ALPs have masses that are proportional to the scale of explicit breaking and are therefore naturally much lighter than f. The approximate shift symmetry enjoyed by ALPs requires that they couple to other fields via derivatives, and the strengths of these couplings are suppressed by powers of $\frac{1}{f}$. Consequently, ALPs are ideal mediators between dark matter (DM) and SM particles through what has been dubbed the "axion portal" [13–22].

If SM fermions are also charged under the PQ symmetry, the ALP couples derivatively to the corresponding fermion current, and the phenomenology resembles that of a spin-0 particle with Yukawa-like couplings [15]. By contrast, when only beyond-SM fermions are charged under both the PQ symmetry and SM gauge interactions, the predominant ALP coupling to the SM is to pairs of SM gauge bosons. Recent studies have comprehensively assessed the status of axion-portal DM with predominant couplings to gluons [20,21].

We consider instead the scenario in which the ALP interacts with electroweak gauge bosons. Such a limit is interesting because it allows direct annihilation of DM into photon pairs via the ALP mediator, and consequently there exist strong constraints from indirect detection searches for photon lines [23]. These bounds on DM annihilation are complemented by a vast array of direct experimental searches for ALPs, as well as constraints on other particles that we have reinterpreted to apply to ALPs [23–61]. We derive the DM-motivated parameter space for the electroweak axion portal scenario, showing the interplay between indirect detection searches for DM and direct probes of the ALP mediator. A similarly motivated study, albeit one with very different phenomenology, was recently performed for the case where the ALP couples to leptons in addition to photons [22].

Since the relative sizes of the DM-ALP and SM-ALP couplings are *a priori* undetermined, we propose simple benchmarks with SM gauge anomaly coefficients that are either "small" or "large" and that parametrize different well-motivated limits of the parameter space. We find that current constraints already set meaningful bounds on parameters predicting the DM abundance through thermal freeze-out and that improvements of one to two orders of magnitude in sensitivity to ALP couplings would test a substantial fraction of the remaining parameter space. Our benchmarks therefore provide insight into gaps in current coverage of ALP-mediated DM and motivation for

future searches. We additionally consider the parameter space for freeze-in, identifying the relevant processes driving the abundance in each parameter regime. ALP-portal freeze-in DM has also been recently considered in Refs. [20,21,62,63].

We first present the ALP effective field theory and benchmarks in Sec. II. We then study the predictions for thermal DM in Sec. III, presenting interplay with indirect detection constraints in Sec. IV and direct searches for ALPs in Sec. V. Finally, we consider the motivated parameter space of and detection prospects for freeze-in DM in Sec. VI.

II. MEDIATOR EFFECTIVE FIELD THEORY

We consider an effecztive field theory (EFT) in which DM, χ, is a Dirac fermion with unit charge under an anomalous PQ symmetry. If the DM is sufficiently weakly coupled to the scalar(s) responsible for breaking the PQ symmetry, only two states remain in the EFT: DM and the ALP, a. We take the ALP mass to be a free parameter: its mass could originate either from explicit PQ breaking or from a confining gauge interaction in the hidden sector. In any case, M_a is technically natural. The validity of the EFT requires that M_χ and M_a be well below the radial-mode masses of the symmetry-breaking sector, which are comparable to the ALP decay constant, f.

Since the ALP is a pseudo-Goldstone boson, its interactions respect a shift symmetry $a \to a + \theta f$ for any constant θ. The EFT should therefore contain only derivative couplings of the ALP to DM, which are automatically invariant under this shift symmetry. In the simplest UV completions, such as the one we present in Appendix C, the leading-order ALP-DM Lagrangian is

$$\mathcal{L}_{\mathrm{DM-ALP}} = -\frac{1}{2f}\partial_\mu a \bar{\chi}\gamma^\mu\gamma^5\chi. \tag{1}$$

We are interested in DM annihilation into ALPs, $\chi\bar{\chi} \to aa$, whose matrix element is $\mathcal{O}(\frac{1}{f^2})$. Thus, we might expect additional higher-order EFT operators to contribute, such as $a\partial_\mu a\bar{\chi}\gamma^\mu\chi$. In the UV completion in Appendix C, such operators are absent and so we take Eq. (1) as our full ALP-DM Lagrangian.

It is common in the literature to perform field redefinitions (sometimes called "applying the equations of motion") to convert the derivative coupling into a polynomial coupling of ALPs to DM, although this leads to more interaction terms as we discuss in Appendix C. We therefore find it simpler to work with the single derivative coupling in Eq. (1). More complicated UV completions could introduce additional operators, but we expect that these would have only modest effects on our results. We neglect additional corrections from PQ-breaking proportional to M_a since we expect these are small in the pseudo-Goldstone limit, although these corrections could become important for $M_a \sim f$.

In addition to its coupling to DM, the ALP can also couple to SM fields. We consider the scenario in which SM fields are neutral under the PQ symmetry; in this case, the ALP does not couple directly to SM fermions. The ALP can, however, couple to SM gauge bosons via intermediate states with masses $\sim f$ carrying both PQ and SM charges. At low energies, the Lagrangian for the ALP-SM couplings is

$$\mathcal{L}_{\mathrm{gauge-ALP}} = -\frac{C_{aB}}{4}aB_{\mu\nu}\tilde{B}^{\mu\nu} - \frac{C_{aW}}{4}aW^b_{\mu\nu}\tilde{W}^{b\mu\nu} - \frac{C_{ag}}{4}aG^b_{\mu\nu}\tilde{G}^{b\mu\nu}, \tag{2}$$

where B, W, and G are respectively the hypercharge, $\mathrm{SU}(2)_L$, and QCD gauge fields, and $\tilde{V}^{\mu\nu} = \frac{1}{2}\epsilon^{\mu\nu\rho\sigma}V_{\rho\sigma}$ is the dual gauge boson field strength. The constants C_{aV} depend on the charges of the heavy states to SM gauge bosons V and are proportional to $\frac{\alpha_V}{\pi f}$.

In our study, we focus on the couplings between ALPs and electroweak gauge bosons. After electroweak symmetry breaking, the physically relevant couplings are those to the photon, $W^\pm$, and Z bosons. The Lagrangian is

$$\mathcal{L}_{\mathrm{gauge-ALP}} = -\frac{g_{a\gamma\gamma}}{4}aF_{\mu\nu}\tilde{F}^{\mu\nu} - \frac{g_{aZZ}}{4}aZ_{\mu\nu}\tilde{Z}^{\mu\nu} - \frac{g_{a\gamma Z}}{2}aF_{\mu\nu}\tilde{Z}^{\mu\nu} - \frac{g_{aWW}}{2}aW^+_{\mu\nu}\tilde{W}^{-\mu\nu}, \tag{3}$$

where

$$g_{a\gamma\gamma} = C_{aB}\cos^2\theta_{\mathrm{W}} + C_{aW}\sin^2\theta_{\mathrm{W}}, \tag{4}$$

$$g_{aZZ} = C_{aB}\sin^2\theta_{\mathrm{W}} + C_{aW}\cos^2\theta_{\mathrm{W}}, \tag{5}$$

$$g_{a\gamma Z} = (C_{aW} - C_{aB})\sin\theta_{\mathrm{W}}\cos\theta_{\mathrm{W}}, \tag{6}$$

$$g_{aWW} = C_{aW}, \tag{7}$$

and θ_{W} is the weak mixing angle.

Additional couplings to SM fermions are generated by renormalization group (RG) evolution from the cutoff of the EFT ($\lesssim 4\pi f$) down to phenomenologically relevant scales [64,65]. These additional couplings are UV sensitive and model dependent, and we generally do not include them in our calculations. We find that the electroweak ALP portal predicts relatively low EFT cutoffs, which makes RG-induced couplings less relevant than they otherwise might be. However, in our analysis, we highlight where these RG-induced couplings could have a significant impact on the phenomenology.

A. Specific model benchmarks

The undetermined coefficients C_{aV} in Eq. (2) do not allow a direct correspondence between the DM and gauge couplings of the ALP. However, it is possible to relate the couplings in specific UV completions. For example, consider a Dirac fermion Ψ that is a hypercharge singlet but vectorlike under $\mathrm{SU}(2)_L$ with representation R. Ψ has an axial charge under the PQ symmetry such that it acquires its mass through the mechanism that spontaneously breaks $\mathrm{U}(1)_{\mathrm{PQ}}$. In this case, Ψ acquires a coupling to the ALP of (upon applying equations of motion[1])

$$\mathcal{L}_{\Psi a} = \frac{iM_\Psi a}{f}\bar{\Psi}\gamma^5\Psi. \tag{8}$$

We can then calculate the ALP coupling to $\mathrm{SU}(2)_L$ gauge bosons explicitly, obtaining

$$C_{aW} = \frac{\alpha_W T(R)}{\pi f}, \tag{9}$$

where $T(R)$ is the index of the representation R defined by $\mathrm{Tr}(T_R^a T_R^b) \equiv T(R)\delta^{ab}$, and T_R^a are generators for the representation R. The coupling in Eq. (9) can also be obtained using anomalous field redefinitions on the corresponding θ-parameter. Note that $T(\square) = \frac{1}{2}$ for the fundamental representation, and this is the minimum that a field Ψ can contribute to the ALP coupling to $\mathrm{SU}(2)_L$ gauge bosons.

Similarly, we can consider a benchmark where Ψ is instead an $\mathrm{SU}(2)_L$ singlet with hypercharge Q_Y. In this case,

$$C_{aB} = \frac{\alpha_Y Q_Y^2}{\pi f}. \tag{10}$$

This motivates us to consider the following benchmarks in our study:

$$(1):\ C_{aW} = 0, \qquad C_{aB} = \frac{\alpha_Y Q_Y^2}{\pi f}, \tag{11}$$

$$(2):\ C_{aW} = \frac{\alpha_W T(R)}{\pi f}, \qquad C_{aB} = 0 \tag{12}$$

for various choices of R and Q_Y.

We are interested in a "small coupling" and "large coupling" for each of the hypercharge and SU(2) couplings. We choose $Q_Y = 1$ and $T(R) = 1/2$ for the small-coupling benchmarks and $Q_Y = 10$ and $T(R) = 110$ for the large-coupling benchmarks.

[1]This is valid here because, in our analysis, only terms in the matrix element that are linear in the gauge boson coupling are relevant.

III. THERMAL DARK MATTER ABUNDANCE

The ALP mediates production and annihilation of DM to/from SM states. Which of these two processes is most relevant depends on the DM-ALP mass hierarchy. In this section, we focus on thermal DM scenarios and we return to freeze-in in Sec. VI. There are two regimes relevant for our thermal DM study:

(1) *Secluded Scenario.* In this scenario, $M_a < M_\chi$ and the dominant annihilation mode is typically $\chi\bar{\chi} \to aa$. The relic abundance is determined entirely by the ALP-DM coupling, which scales like $\frac{1}{f}$. However, the ALP-SM couplings also scale like $\frac{1}{f}$ and the two are therefore related in a manner that is parametrically different from renormalizable portals. Thus, the electroweak ALP portal provides relatively concrete predictions of phenomenological signatures even in the secluded scenario.

(2) *Annihilation to SM.* In this scenario, $M_a \geq M_\chi$ and DM annihilation to ALPs is kinematically forbidden at zero temperature. The dominant annihilation mode is instead $\chi\bar{\chi} \to \gamma\gamma$ (or other SM gauge boson pairs). This regime faces strong constraints from gamma-ray line searches, except when DM annihilates resonantly through the ALP in the early universe but not in the present [23].

(3) *Study Methodology.* We determine the viable parameter space consistent with the observed DM abundance in the following two ways:
   (i) We implement our model into the CalcHEP format[2] [66] using FeynRules [67], and we then use microMEGAs5.2.13 [68] to solve for the DM abundance.
   (ii) We derive and solve a thermally averaged Boltzmann equation for χ that includes all annihilation and decay processes (see Appendixes A and B for more details on our Boltzmann equation and rates, respectively). In thermal scenarios, we can assume that the ALPs are always in equilibrium.

In our Boltzmann equation, the only ALP-SM interaction is with the photon, whereas in microMEGAs we include all SM gauge bosons. We have checked that the predictions for the DM abundance of the two approaches agree to within 20% for most parameters within the regime of mutual validity, $M_a \lesssim M_Z$. We find stronger disagreement very close to resonance, $(M_a - 2M_\chi)/M_a < 0.01$; in this limit, our calculation of the thermally averaged cross section agrees with the prediction in the narrow-width approximation. For most parameter points, we use microMEGAs because it includes all gauge interactions. However, very close to resonance, we use our own solution to the Boltzmann

[2]The calchep model files are available at https://github.com/bshuve/ALP_CalcHEP.

equations, with the on shell cross section modified by appropriate powers of the photon branching fraction when the ALP can decay into multiple gauge bosons.

Because of subtleties in the treatment of annihilation near resonance, we now provide more details on the relevant rates in that region of parameter space.

A. DM annihilation near resonance

When $M_a \approx 2M_\chi$, there is a resonant enhancement of the DM annihilation cross section. We must evaluate the thermally averaged cross section,

$$\langle \sigma_{\bar{\chi}\chi\to\gamma\gamma} v\rangle = \frac{g_{a\gamma\gamma}^2}{1024\pi M_\chi^2 f^2 T K_2(M_\chi/T)^2} \times \int_{4M_\chi^2}^{\infty} \frac{s^{7/2}\sqrt{1-4M_\chi^2/s}}{(s-M_a^2)^2+M_a^2\Gamma_a^2} K_1(\sqrt{s}/T) ds, \quad (13)$$

in this limit, where K_1 and K_2 are modified spherical Bessel functions of the second kind and s is the Mandelstam variable. The correct treatment of resonance depends on whether DM is heavier or lighter than $M_a/2$; in the latter case, the parametric dependence of the annihilation cross section is also sensitive to whether the ALP predominantly decays into DM or SM gauge bosons. To simplify our analysis, in this section we restrict our consideration to $M_\chi \lesssim 45$ GeV, in which case the only open annihilation modes to the SM in the resonance region are to pairs of photons.

Those who are interested only in our quantitative results can skip to the summary of our findings at the end of this section. We now provide details of our calculation of resonant annihilation in various limits.

Above but close to resonance, $M_a > 2M_\chi$, the ALP can decay into both DM and photons. In the narrow-width approximation, the ALP propagator can be approximated with a Dirac delta function,

$$\frac{1}{(s-M_a^2)^2+M_a^2\Gamma_a^2} \to \frac{\pi}{M_a\Gamma_a}\delta(s-M_a^2), \quad (14)$$

allowing a straightforward evaluation of the thermally averaged cross section

$$\langle \sigma_{\bar{\chi}\chi\to\gamma\gamma} v\rangle = \frac{g_{a\gamma\gamma}^2 M_a^6 \sqrt{1-4M_\chi^2/M_a^2} K_1(M_a/T)}{1024 f^2 M_\chi^2 \Gamma_a T K_2(M_\chi/T)^2}. \quad (15)$$

Because we generally expect the ALP coupling to DM to be larger than its coupling to photons, the decay width is dominated by the decay to DM even for M_a quite close to $2M_\chi$. In this case, we can approximate $\Gamma_a \approx \Gamma(a\to\chi\bar{\chi})$, where

$$\Gamma(a\to\bar{\chi}\chi) \approx \frac{M_a M_\chi^2}{8\pi f^2}\sqrt{1-\frac{4M_\chi^2}{M_a^2}}. \quad (16)$$

Substituting this into the thermally averaged cross section gives

$$\langle \sigma_{\bar{\chi}\chi\to\gamma\gamma} v\rangle \approx \frac{\pi g_{a\gamma\gamma}^2 M_a^5 K_1(M_a/T)}{128 M_\chi^4 T K_2(M_\chi/T)^2}. \quad (17)$$

Note that this rate depends only on the ALP coupling to SM fields and not on the DM coupling. It does, however, implicitly assume that the DM coupling is much larger than $g_{a\gamma\gamma}$ for the decay to DM to dominate. The parametric dependence of the cross section in Eq. (17) agrees with that found in Ref. [23]. In particular, for $M_a \approx 2M_\chi$ and freeze-out occurring at $M_\chi/T \approx 20$, this yields

$$\langle \sigma v\rangle \approx (6\times 10^{-26}\ \mathrm{cm}^3/\mathrm{s})\left(\frac{g_{a\gamma\gamma}}{1.2\times 10^{-5}\ \mathrm{GeV}^{-1}}\right)^2, \quad (18)$$

such that $g_{a\gamma\gamma} \approx 1.2\times 10^{-5}$ GeV^{-1} will give a thermal cross section. This prediction is independent of M_χ.

The above estimate, however, does not apply very near to resonance. $\Gamma(a\to\bar{\chi}\chi)$ goes to zero in the limit $M_a \to 2M_\chi$ while the decay width to photons,

$$\Gamma(a\to\gamma\gamma) = \frac{g_{a\gamma\gamma}^2 M_a^3}{64\pi}, \quad (19)$$

does not. Therefore, sufficiently close to resonance, the ALP decay to photons dominates over the decay to DM. When this happens (while continuing to use the delta-function approximation of the ALP propagator), the thermally averaged cross section depends only on the coupling to DM

$$\langle \sigma_{\bar{\chi}\chi\to\gamma\gamma} v\rangle \approx \frac{\pi M_a^3 K_1(M_a/T)\sqrt{1-4M_\chi^2/M_a^2}}{16 M_\chi^2 f^2 T K_2(M_\chi/T)^2}. \quad (20)$$

Note that the cross section appears to vanish in the $M_a \to 2M_\chi$ limit!.

The reason this appears to go to zero on resonance is that the delta function only has support at $s = M_a^2$, but the phase space also vanishes in this limit. We must therefore be more careful with the thermal average integral, treating Γ_a as small but nonvanishing when we take $M_a \to 2M_\chi$. To illustrate our approach, we evaluate the integral exactly on resonance, $M_a = 2M_\chi$. If the width is narrow, we can approximate $s \approx M_a^2 = 4M_\chi^2$ everywhere *except* in two places: in the denominator and in the phase space factor in the numerator

$$\langle \sigma_{\bar{\chi}\chi\to\gamma\gamma} v\rangle_{\mathrm{res}} = \frac{g_{a\gamma\gamma}^2 M_a^5 K_1(M_a/T)}{256\pi f^2 T K_2(\frac{M_a}{2T})^2} \times \int_{M_a^2}^{\infty} \frac{\sqrt{1-M_a^2/s}}{(s-M_a^2)^2+M_a^2\Gamma_a^2} ds. \quad (21)$$

The integral can be evaluated analytically assuming $\Gamma_a \ll M_a$, in which case

$$\langle \sigma_{\bar{\chi}\chi\to\gamma\gamma} v\rangle_{\rm res} \approx \frac{g_{a\gamma\gamma}^2 M_a^5 K_1(M_a/T)}{256 f^2 T K_2(\frac{M_a}{2T})^2} \frac{1}{\sqrt{2M_a^3\Gamma_a}}. \quad (22)$$

As expected, we find the thermally averaged cross section is nonzero even exactly on resonance.

The nonanalytic dependence on the width gives rise to interesting scaling behavior. For example, substituting the ALP decay width to SM photons gives

$$\langle \sigma_{\bar{\chi}\chi\to\gamma\gamma} v\rangle_{\rm res} \approx \sqrt{\frac{\pi}{2}} \frac{g_{a\gamma\gamma} M_a^2 K_1(M_a/T)}{32 f^2 T K_2(\frac{M_a}{2T})^2}, \quad (23)$$

with a linear dependence on the coupling to photons. Defining $x \equiv \frac{M_\chi}{T} = \frac{M_a}{2T}$, which is typically ∼15–20 for thermal freeze-out, we have

$$\langle \sigma_{\bar{\chi}\chi\to\gamma\gamma} v\rangle_{\rm res} \approx \sqrt{\frac{\pi}{2}} \frac{g_{a\gamma\gamma} M_a x K_1(2x)}{16 f^2 K_2(x)^2}. \quad (24)$$

The linear dependence of the cross section on the coupling has interesting implications. For example, our benchmarks in Sec. II A all have $g_{a\gamma\gamma} \sim \frac{1}{f}$. If we want Eq. (23) to give a fixed thermal cross section to achieve the correct relic abundance, we predict a $\frac{1}{f} \propto \frac{1}{M_a^{1/3}}$ scaling for the thermal benchmark.

When $M_a < 2M_\chi$, the resonance condition cannot be exactly satisfied. If $0 < (4M_\chi^2 - M_a^2) \lesssim M_a\Gamma_a$, the annihilation still occurs close enough to resonance that the above calculation applies in an approximate manner. Away from resonance, the cross section plummets rapidly for fixed couplings and the results are the same as those for annihilation through an off shell ALP.

For $M_a \gtrsim 90$ GeV, additional decay modes open to γZ, ZZ, and WW depending on the ALP mass. This leads to a rescaling of the resonant annihilation cross section. Above resonance, with the ALP predominantly decaying to DM, Γ_a is unchanged by an increase in the effective coupling to the SM while the numerator increases due to more DM annihilation modes. Consequently, $\langle \sigma_{\bar{\chi}\chi} v\rangle$ increases by $\mathrm{BF}(a\to\gamma\gamma)^{-1}$ relative to the photon-only scenario. Above—but sufficiently close to resonance that the ALP decays predominantly to SM states—$\langle \sigma_{\bar{\chi}\chi} v\rangle$ is independent of the SM coupling and is therefore independent of $\mathrm{BF}(a\to\gamma\gamma)$. Finally, exactly on resonance we have that the cross section scales as $\mathrm{BF}(a\to\gamma\gamma)^{-1/2}$ due to the enhancement from additional SM annihilation modes being partially cancelled by the factor of $\Gamma_a^{-1/2}$ in the denominator.

Summary of findings:

(i) When $M_a \lesssim 2M_\chi$ but very close to resonance, the correct thermally averaged cross section is given by Eq. (23);

(ii) When $M_a > 2M_\chi$, we determine for the given ALP mass whether the decay is predominantly to DM or to photons. If it decays predominantly to DM, then Eq. (17) applies. Otherwise, Eq. (20) applies.

IV. INDIRECT DETECTION CONSTRAINTS ON DARK MATTER

A. Overview

Given that no definitive electromagnetic signals of DM have been discovered to date, indirect detection sets powerful constraints on DM coupled to photons. In this section, we apply photon-line searches and CMB ionization bounds to the electroweak ALP portal, comparing constraints to parameters predicted by the DM thermal relic abundance.

The process $\bar{\chi}\chi \to \gamma\gamma$ in the DM halo today leads to a monochromatic photon-line feature with $E_\gamma = M_\chi$. This is an s-wave scattering process at low momentum, and consequently it is unsuppressed at late times. In contrast, the process $\bar{\chi}\chi \to aa$, when kinematically allowed, occurs in the p-wave; even in the secluded scenario when this process determines the DM abundance, the direct annihilation to photons provides the strongest bounds today.

In the limit $v \to 0$, the thermally averaged diphoton cross section is

$$\langle \sigma_{\bar{\chi}\chi\to\gamma\gamma} v\rangle_{v\to 0} = \frac{g_{a\gamma\gamma}^2 M_\chi^6}{4\pi f^2} \frac{1}{(4M_\chi^2 - M_a^2)^2 + M_a^2\Gamma_a^2}. \quad (25)$$

We now discuss constraints on this annihilation rate from astrophysical and cosmological observables.

Photon-line searches. We apply constraints on $\langle \sigma_{\bar{\chi}\chi\to\gamma\gamma} v\rangle$ from galactic photon-line searches with Fermi-LAT [69,70], MAGIC [71], H.E.S.S [72], EGRET, and COMPTEL [73,74].[3] Because constraints from gamma-ray lines reach subthermal cross sections, indirect detection bounds are relevant even when $\bar{\chi}\chi \to \gamma\gamma$ is a subdominant DM annihilation mode. These constraints are sensitive to the DM profile in the Milky Way, especially in the Galactic center region. For concreteness, we present bounds assuming a Navarro-Frenk-White profile [76]. Due to uncertain systematic uncertainties affecting modelling of the photon continuum in the EGRET dataset, we present both conservative [74] and optimistic limits [73] from EGRET. Ref. [74] also includes constraints from COMPTEL.

[3]Current searches have low sensitivity to gamma rays in the MeV-GeV energy range. These can be improved with future telescopes such as e-ASTROGAM [75].

There are, in principle, constraints on continuum gamma-ray production, either from DM annihilation to ALP pairs or, at high masses, to W and Z bosons. Since these bounds are weaker than those from photon lines (see, e.g., [77]) we do not include them in our analysis.

Cosmic microwave background. DM annihilation to photons during recombination leads to distortions in the cosmic microwave background (CMB) that are incompatible with Planck data [78]. We apply the constraints from Ref. [79] to our model.

Along with the above constraints, we impose two model-dependent consistency conditions. For the first, we require that there exists a perturbative UV completion of the EFT such that DM annihilation occurs within the regime of validity of the EFT. Perturbative unitarity bounds the radial mode of the PQ-breaking scalar to have a mass that is $\lesssim\sqrt{4\pi}f$, where we take couplings $\lesssim\sqrt{4\pi}$ for the theory to remain unitary once group theory multiplicity factors are taken into account. Furthermore, the typical momentum flowing through the ALP is $2M_\chi$ during late-time annihilation, which together leads to the consistency bound of $M_\chi \leq \sqrt{\pi}f$. For the ALP to be the dominant mediator between DM and the SM, we additionally require $M_a \leq \sqrt{4\pi}f$. Note that these bounds on the validity of the EFT are approximate and depend to some extent on the UV completion, with a dedicated study needed to determine the precise limit of validity.

The second consistency condition comes from the fact that we have not observed direct production of the charged fermions, Ψ, which mediate interactions between the ALP and SM gauge bosons. Since the EFT depends only on the sum of the anomaly coefficients of all such particles, it is not possible to directly impose such a bound. For the purposes of comparison, we derive approximate constraints assuming the existence of a single such fermion with a mass equal to $\sqrt{4\pi}f$.

For the hypercharge-coupled case and $Q_Y = 1$, the Large Electron-Positron collider (LEP) constrains new charged fermion masses to be above 90 GeV in minimal models [80]. For $Q_Y = 10$, the new charged states form bound states that annihilate to photon pairs, and they can be constrained from an ATLAS search for diphoton resonances [56,81]; this leads to a stronger constraint than searches for heavy stable charged particles. Using Ref. [81], we estimate a bound of 3 TeV on heavy fermion masses, corresponding to a bound-state mass of 6 TeV.

When the ALP couples predominantly to $\mathrm{SU}(2)_L$ gauge bosons, the fundamental representation has fermions of charge $\pm\frac{1}{2}$ for which a CMS search for fractionally charged particles constrains masses to be above 600 GeV [82]. The value $T(R) = 110$ corresponds to an eleven-plet with maximum charge of 5, and so the ATLAS diphoton search constrains the mass to be above approximately 2.2 TeV [56,81].

We reiterate that each of these constraints on new charged states is highly model dependent, and in many cases the same ALP-gauge coupling can be approximately achieved with drastically changed bounds on the charged state(s).

Finally, we note that direct-detection constraints on DM are expected to be negligible for DM through the electroweak axion portal [83].

B. Secluded scenario results

In Fig. 1, we present our results for the secluded scenario, $M_a < M_\chi$. When $M_a \ll M_\chi$, we find that the results are independent of the precise M_a/M_χ ratio, and so we choose $M_a/M_\chi = 1/3$ as a benchmark for our study. In this limit, the DM annihilation cross section to ALPs scales like $\frac{M_\chi^2}{f^4}$, leading to a prediction of $\frac{1}{f} \sim \frac{1}{\sqrt{M_\chi}}$ scaling for the thermal abundance. This scaling can be seen in our numerical results for the DM abundance curve, although there are deviations in this scaling for masses in the 1–10 GeV range due to rapid changes in g_* with temperature during freeze-out for these masses (this effect also shows up in other scenarios). Additionally, when large couplings exist to W and Z bosons (for example, with $\mathrm{SU}(2)_L$ coupling and $T(R) = 110$), the annihilation rate to W and Z bosons can be large enough to induce deviations in this scaling in the vicinity of the diboson mass threshold.

In the "small" gauge coupling limit ($Q_Y = 1$ or $T(R) = 1/2$), indirect detection constraints are not yet powerful enough to constrain much of the DM parameter space. Improvements to photon-line constraints could meaningfully probe much of the remaining parameter space. By contrast, for "large" gauge couplings ($Q_Y = 10$ or $T(R) = 110$), the predicted DM coupling drops slightly but the constraints get much stronger. Consequently, the DM parameter space is almost completely excluded, except for a small corner in the vicinity of $M_\chi = 10$ TeV. We also see that the results are similar for the hypercharge- and $\mathrm{SU}(2)_L$-coupled scenarios. For the rest of this section, we only show results for the hypercharge-coupled scenario although the outcome is similar for both.

For nearly the whole parameter space, f cannot be much above the weak scale and deviations from the predicted EFT behavior are possible. This requires a dedicated analysis that is beyond the scope of our work, but our results show that care must be taken in extrapolating results of the electroweak axion portal to energies at a TeV or above.

C. Annihilation to SM far from resonance

In the $M_a \geq M_\chi$ regime, DM annihilation could proceed below, on, or above resonance. We defer to the next section an exploration of the resonance region, and so we now

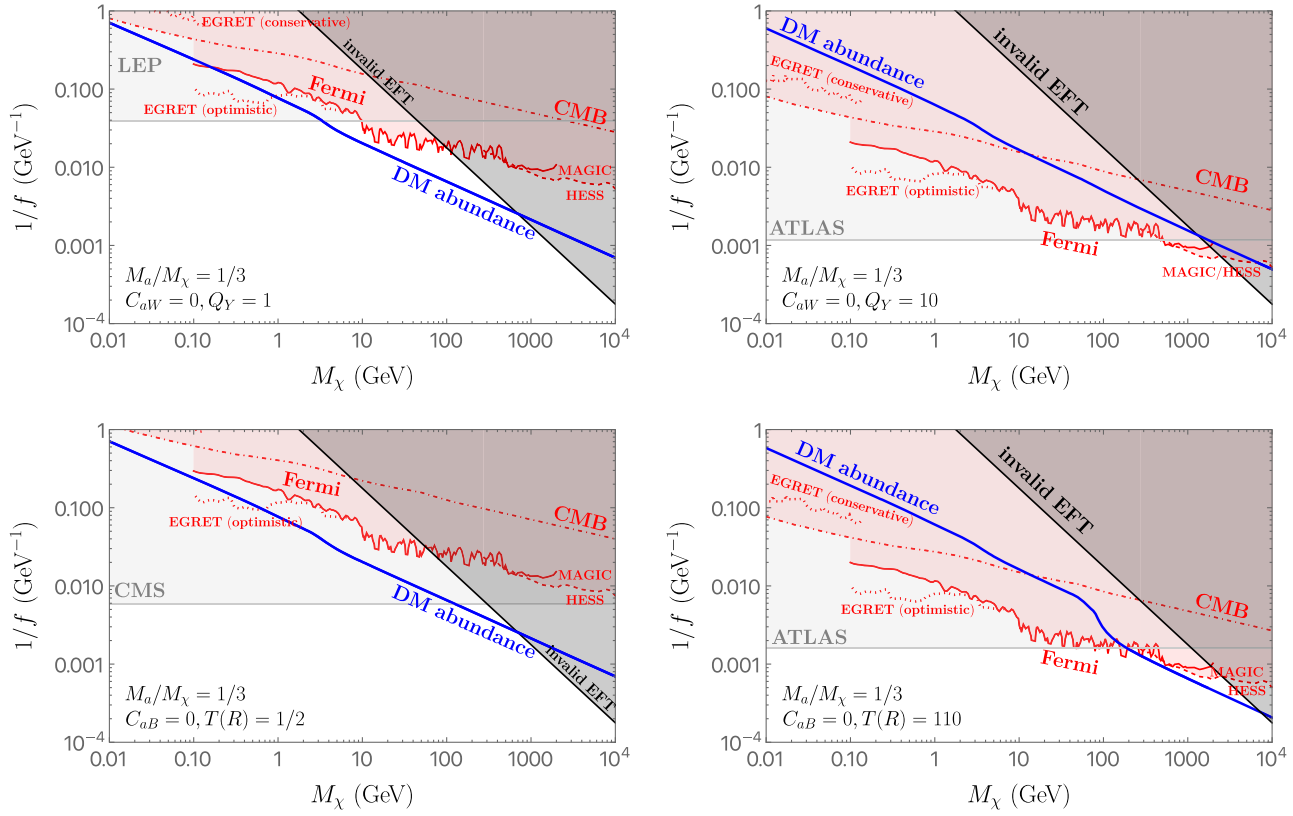

FIG. 1. DM thermal-relic parameters along with indirect-detection constraints for the secluded scenario ($M_a/M_\chi = 1/3$). The top row shows results for the hypercharge-coupled scenario with (left) $Q_Y = 1$; (right) $Q_Y = 10$; the bottom row shows results for the $SU(2)_L$-coupled scenario with (left) $T(R) = 1/2$ and (right) $T(R) = 110$. Constraints are from: (solid) Fermi-LAT; (dashed) MAGIC and H.E.S.S.; (dot-dashed) CMB; (dotted) EGRET. The region shaded in dark gray is inconsistent with our EFT framework, while the light gray shaded region indicates model-dependent collider constraints on the charged states in the minimal UV completion of the model.

choose mass ratios far below ($M_\chi/M_a = 1.5$) and above ($M_\chi/M_a = 3$) resonance. We show the results of this study in Fig. 2.

In this limit, the s-wave $\bar{\chi}\chi \to \gamma\gamma$ process occurs at comparable rates during DM freeze-out and at later times. Therefore, the relative relationship between the DM abundance curve and indirect detection constraints is largely independent of couplings. As expected, we see that indirect detection searches consistently rule out the thermal relic parameters.

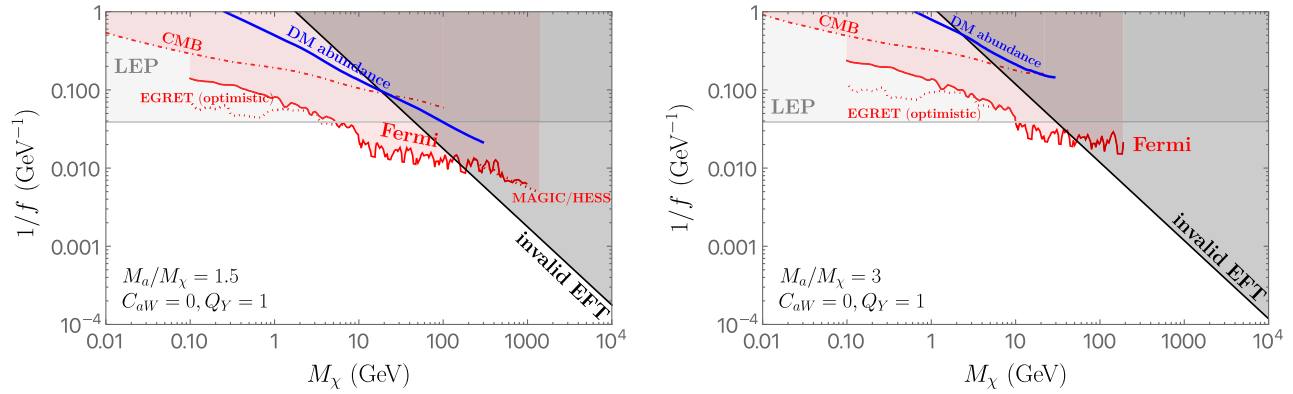

FIG. 2. DM thermal-relic parameters along with indirect-detection constraints for $M_a > M_\chi$ far from resonance. In the left (right) figure, we show results for DM annihilating below (above) resonance, with $M_a/M_\chi = 1.5$ ($M_a/M_\chi = 3$). All curves and shaded regions are the same as in Fig. 1; we do not include MAGIC or H.E.S.S. results for the right plot because they lie entirely in the region of invalid EFT.

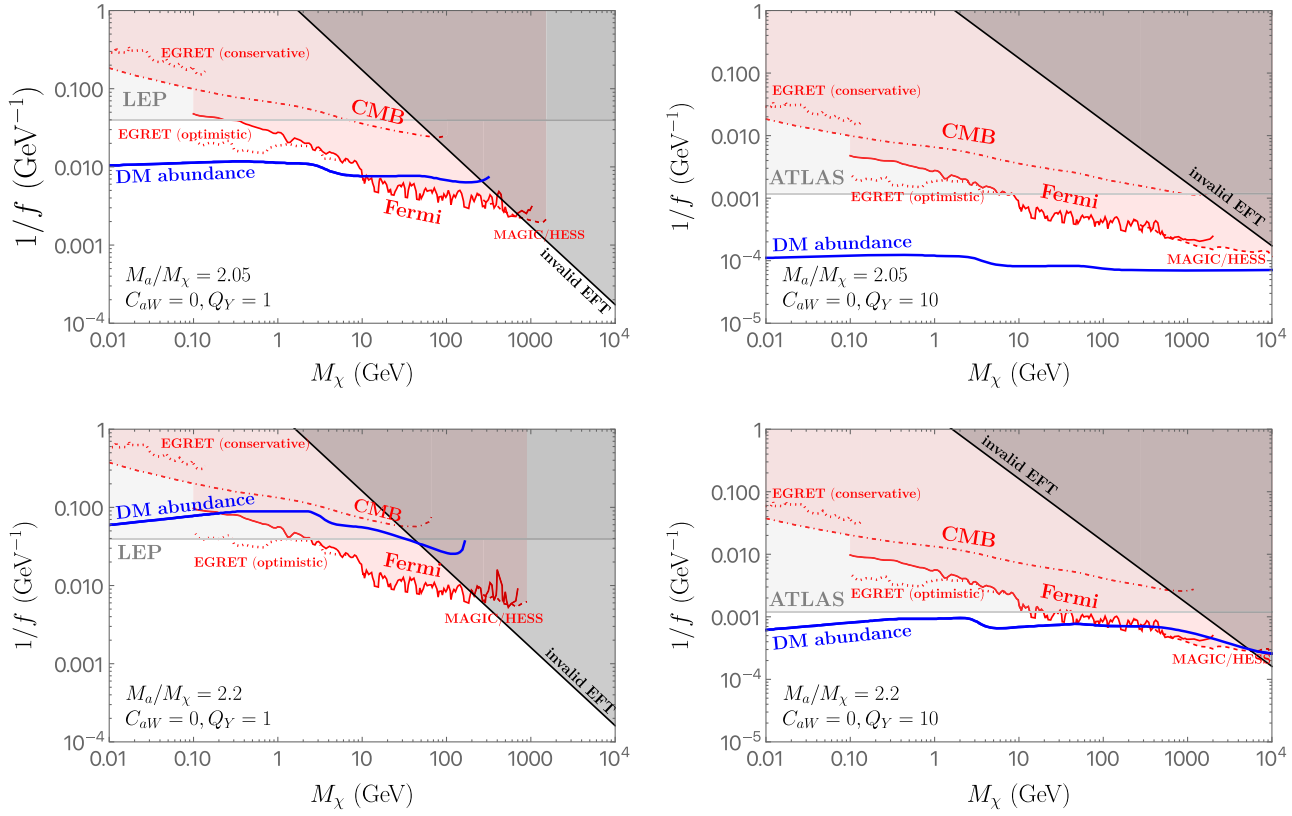

FIG. 3. DM thermal-relic parameters along with indirect-detection constraints in the resonance region. In the top ($M_a/M_\chi = 2.05$) and bottom ($M_a/M_\chi = 2.2$) rows, we show results for DM annihilation near resonance; the left (right) plots assume $Q_Y = 1$ ($Q_Y = 10$). All curves and shaded regions are the same as in Fig. 1.

D. Annihilation on resonance

Perhaps the most interesting scenario is within the resonance region, $M_a/M_\chi \approx 2$. When the DM-ALP mass splitting is such that there is a resonant enhancement of annihilation to photons in the early Universe but *not* at later times, it is possible to achieve the observed relic abundance consistent with indirect detection bounds; see Fig. 3. For the parameter space we show, DM annihilation occurs close to resonance but in the regime where ALP decays to DM still dominate. In this case, Eq. (17) holds and $\langle\sigma v\rangle \sim g_{a\gamma\gamma}^2 \sim \frac{\alpha_Y^2 Q_Y^4}{\pi^2 f^2}$. Since the thermal relic cross section is a constant, this gives a thermal scaling relation of $\frac{1}{f} \sim \frac{1}{Q_Y^2}$, consistent with the DM abundance curves in Fig. 3.

The smaller predicted DM-ALP coupling lessens the constraints from indirect detection. Therefore, when annihilating near resonance, we see the opposite of our findings from the secluded scenario: in that case, larger ALP-gauge coupling led to strong indirect detection constraints, whereas here the larger gauge coupling reduces tension between indirect-detection bounds and the DM abundance.

For most DM masses, the thermal relic curve is approximately independent of M_χ, again in agreement with the prediction of Eq. (17). This behavior breaks down for parameters near the boundary of EFT validity, in which case the ALP coupling is relatively strong and the narrow-width approximation no longer applies.

V. ACCELERATOR AND COLLIDER SEARCHES FOR ALPS

In recent years, there have been a plethora of new searches for ALPs in accelerators and colliders, motivated by the ubiquity of ALPs in new physics models. Our earlier results showed that the electroweak axion portal gives rise to concrete predictions for the ALP mass and coupling consistent to achieve the correct DM abundance, and we can therefore compare the DM-motivated parameter space to constraints from direct searches for ALPs.

Search strategies for ALPs depend on whether or not the ALP decays predominantly to SM gauge bosons or DM. In the former case, the ALP invariant mass can be reconstructed from its decay products, whereas in the latter case its existence must be inferred through missing momentum. Except just above resonance, we find that the ALPs decay visibly for most of the viable parameter space that is not already ruled out by indirect detection, and we therefore focus on this scenario. Previous studies have been done to investigate the phenomenology of invisibly decaying ALPs [23,42,84], with connections to DM explored in the resonant annihilation regime [23].

We find that terrestrial ALP searches already meaningfully constrain ALP portal DM, and improvements in sensitivity to the coupling by an additional order of magnitude or two would allow for a discovery or exclusion of the majority of the parameter space. Just as for indirect detection, much of the "large" gauge coupling parameter space is already ruled out, while the "small" gauge coupling parameter space is more open.

We first provide an overview of the main categories of ALP searches and then present our results in Sec. V B.

A. Overview of searches

Searches for ALPs fall into one of the following four broad categories:

(i) *Hadron colliders:* the large center of mass energy of hadron colliders means that these experiment provide the dominant constraints on high-mass ALPs. However, the large quantity of Z and Higgs bosons produced at hadron colliders also provides sensitivity to rare but spectacular decays into lower-mass ALPs. Heavy-ion collisions at the LHC further provide a unique window into ALPs in the 10–100 GeV mass range.

(ii) *Electron-positron colliders:* ALPs can be directly produced in e^+e^- collisions through a virtual γ/Z or an on shell Z boson. Precision studies of rare Z boson decay modes at LEP provide excellent sensitivity to ALPs produced in their decays.

(iii) *Fixed-target and beam-dump experiments:* these experiments provide limited kinematic reach but extremely high luminosities, which provide sensitivity to very feebly coupled ALPs at lower masses.

(iv) *Flavor experiments:* when ALPs couple to W bosons, they can be emitted in flavor-changing neutral current (FCNC) decays [42]. The rare nature of these decays in the SM means that flavor experiments have excellent sensitivity to ALPs with masses below the B-meson mass.

We show the constraints on the hypercharge- and $SU(2)_L$-coupled scenarios in Fig. 4.

We now provide more details on each of the searches. When needed, signal simulations are done using MadGraph5_aMC@NLO [85] with a UFO [86] model file that we created using FeynRules [67].

Hadron colliders: there are dedicated searches by ATLAS [60] and CMS/TOTEM [61] for photon-coupled ALPs in proton-proton collisions at the LHC. In these searches, the ALP is created in light-by-light scattering with one or both of the protons remaining intact and being tagged in the forward region. These searches collectively span the ALP mass range from 150 GeV to 2 TeV. Because the results provided by the experimental collaborations are already expressed in terms of the ALP-photon coupling, we can directly map these constraints to our coupling benchmarks.

ALPs are also constrained by inclusive diboson searches. ATLAS has performed an inclusive search for $\gamma\gamma$ resonances [56], reporting their results in terms of the fiducial cross section times diphoton branching fraction. We simulate inclusive production of the ALP at parton level, for which the dominant channels are vector-boson-fusion (VBF), $pp \to a + jj$ and associated production, $pp \to a + V$ for some gauge boson V. We then decay the ALP into photons and use our simulated events to calculate the fiducial cross section, which ranges from 60–80% of the total cross section. Finally, we translate the ATLAS bounds onto the ALP coupling parameter space for masses 160 GeV to 3 TeV. This sets the strongest bounds on the lower part of the mass range. This is true even for the $SU(2)_L$-coupled scenario because the cross section bounds are stronger than other low-mass diboson searches even after the smaller $\gamma\gamma$ branching fraction is taken into account.

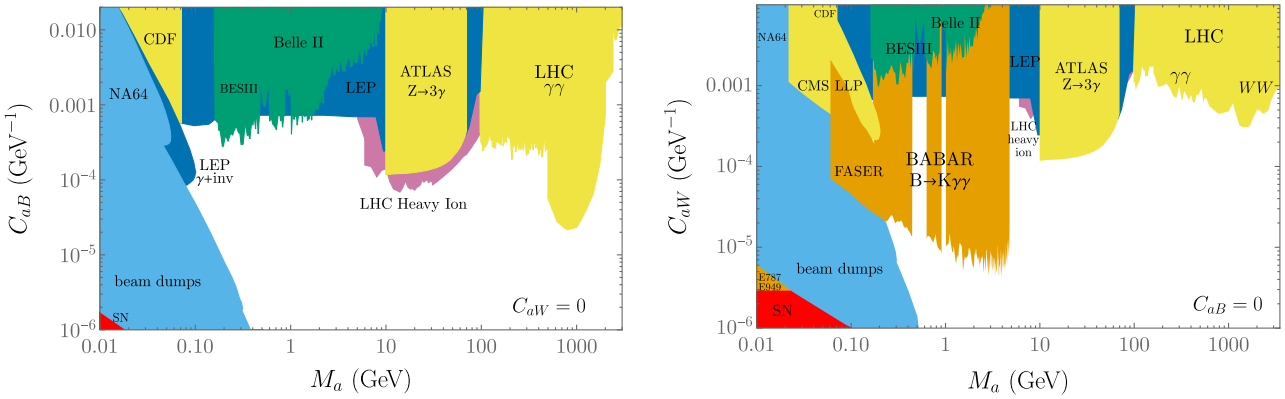

FIG. 4. Bounds from accelerator and collider experiments on (left) the hypercharge-coupled scenario; (right) the $SU(2)_L$-coupled scenario. Bounds from hadron colliders are in yellow; heavy ion collisions in purple; high-energy electron-positron colliders in dark blue; low-energy electron-positron colliders in green; fixed-target and beam-dump experiments in cyan; and flavor experiments in orange. There is also a sliver of parameter space excluded from supernovae constraints depicted in red. Details of each constraint are found in the text.

These limits extend those derived from earlier LHC searches [33–36].

When the ALP couples to $\mathrm{SU}(2)_L$ gauge bosons, the coupling to photons is suppressed by $\tan^2\theta_W$ compared to a hypercharge coupling. This dramatically weakens the bounds from light-by-light scattering. This motivates performing additional dedicated ALP searches in W^+W^-, ZZ, and γZ final states, which have much larger branching fractions than $\gamma\gamma$ in this case. To our knowledge, no such dedicated searches have yet been performed. ATLAS [44] and CMS [55] have set constraints on VBF production of a narrow spin-0 WW resonance, however, and we can directly apply these bounds to the ALP scenario for masses 300 GeV to 4.5 TeV. Since the ALP decay width scales like M_a^3, the width-to-mass ratio increases monotonically with M_a and so we truncate our plots when Γ_a/M_a exceeds 10%, which is approximately the mass resolution of these searches. The WW bound sets the most powerful constraint for the $\mathrm{SU}(2)_L$-coupled scenario for $M_a \gtrsim 1.2$ TeV.

We additionally check an ATLAS search for γZ resonances [87] where the Z decays leptonically. This search is inclusive with respect to production mode but requires that the leptons firing the trigger be part of the reconstructed resonance mass. We include both VBF and associated production in calculating the ALP constraint, but in the case of associated production we conservatively require that the associated gauge boson decay hadronically to avoid the possibility that the event would be triggered on leptons from the associated gauge boson. This removes at most 15% of the associated production events. We then map the cross section bounds to the ALP coupling. We find that the sensitivity of this search is poorer than in other channels.

Hadron colliders can also probe ALPs much lighter than the weak scale. Light-by-light scattering in ultraperipheral heavy-ion collisions is a sensitive probe of lower-mass ALPs (5–100 GeV) due to the much larger electromagnetic field sourced by ions of large atomic number [41]. We apply constraints from ALP searches in heavy-ion collisions at CMS [45] and ATLAS [48].

Additionally, the large rate of Z boson production at hadron colliders allows them to have substantial sensitivity to $Z \to \gamma a, a \to \gamma\gamma$. Depending on the ALP mass, the two photons in its decay can either be resolved separately or merged into a single electromagnetic shower. Bounds on the ALP coupling have been derived from an ATLAS search for $Z \to 3\gamma$ [39,43]—which has sensitivity above 10 GeV—and a CDF search for $Z \to 2\gamma$, where the ALP photons are merged for ALP masses below about 75 MeV [37,43].

Finally, CMS has recently performed a search for long-lived particles decaying in the muon spectrometer [54]. The muon spectrometer is used as a sampling calorimeter, giving the search sensitivity to diphoton decays [59]. Ref. [59] did a detailed reinterpretation of the search and presented results for an ALP coupled to $\mathrm{SU}(2)_L$ bosons for masses below 200 MeV, and we include these bounds. They did not do a similar study for ALPs coupled to hypercharge, although the sensitivity is expected to be reduced because of the shorter lifetime associated with the larger diphoton coupling, along with the smaller production cross section from electroweak boson decays. As this search affects only a narrow sliver of parameter space, we leave a reinterpretation of this search to the hypercharge-coupled case to future work.

Electron-positron colliders: electron-positron colliders currently set strong constraints on ALPs below the Z mass. We apply bounds previously derived from LEP on $Z \to 3\gamma$, and on $Z \to \gamma\gamma$ where the two photons from ALP decay have been merged [26,27,40]. At smaller ALP masses, it can be long-lived and decay to photon pairs outside the detector. In this regime, LEP constraints on Z boson decays to a photon and an invisible particle put bounds on ALP parameters [23,28,30,59].

Searches for ALPs decaying to photons have also been performed at low-energy e^+e^- colliders. Belle II has a search for ALPs below 10 GeV in associated production with a photon, $e^+e^- \to \gamma^* \to \gamma a, a \to \gamma\gamma$ [49]. BESIII has two searches in radiative J/ψ decays, $J/\psi \to \gamma a, a \to \gamma\gamma$, which are sensitive to ALP masses below 3 GeV [57,88].

Fixed-target and beam-dump experiments: the kinematic reach of fixed-target and beam-dump experiments is limited to sub-GeV ALPs, but these experiments have excellent sensitivity to smaller couplings and/or longer lifetimes than other experiments. The strongest constraints from electron beam dumps come from the E137 experiment at SLAC, which places bounds on the photon coupling [23,24]. The NA64 electron fixed-target experiment probed ALP production via the Primakoff process and subsequent displaced decay of the ALP, either reconstructed as a photon pair inside one of the downstream calorimeters or reconstructed as missing energy if the decay occurs outside the detector [51].

For proton beam dumps, the strongest constraints come from the NuCal experiment [25], and the bounds have recently been updated using improved modelling of ALP production and decay for the hypercharge and $\mathrm{SU}(2)_L$ coupling scenarios [58]. There is a sliver of parameter space at low mass/coupling that is not excluded by NuCal. This is covered by supernova bounds; see Sec. VI C for a more detailed discussion of astrophysical and cosmological bounds on ALPs.

The PrimEx experiment used a photon beam to study the π^0 decay width via the Primakoff effect [32], but the diphoton spectrum can also be used to look for ALP decays into photons [47]. We find that PrimEx constraints are subdominant to other bounds for our benchmarks.

Flavor experiments: when ALPs couple to $\mathrm{SU}(2)_L$ gauge bosons, they can be emitted in FCNC decays [42]. The most powerful flavor constraint is set by a *BABAR* search for $B^\pm \to K^\pm a, a \to \gamma\gamma$ [53]. At lower ALP

mass/coupling, there are constraints from E787 and E949 on a $K^{\pm}$ decaying to a $\pi^{\pm}$ plus an invisible particle [89–91] (where the ALP is sufficiently long-lived that it escapes the detector), as well as on ALP production in FCNC decays of strange and B mesons at the LHC, followed by diphoton decay inside of the FASER experiment [92]. Additional constraints on invisible ALP decays from NA62 and KOTO—and diphoton ALP decays from E949 and KTeV—are subdominant to the other bounds [29,31,38,42,46,50,52].

While we have ignored RG-induced couplings because they are typically subdominant, flavor observables are one area where they can have a significant effect on bounds, especially due to the RG-induced ALP-muon coupling. For example, constraints on the decay $B \to K^{(*)}a, a \to \mu^+\mu^-$ [93,94] are several orders of magnitude stronger than $B^{\pm} \to K^{\pm}a, a \to \gamma\gamma$ and can therefore be competitive with the *BABAR* ALP search [42,65]. We do not include model-dependent dimuon constraints and do not anticipate that they dominate for EFT cutoffs at the 100 GeV to 1 TeV scale, but the interplay of these constraints should be checked in any specific UV completion of the ALP EFT.

B. Results

We first present in Fig. 4 the constraints on visibly decaying ALPs for each of the gauge-coupling scenarios, independent of any DM considerations. In many parts of the parameter space, the constraints are comparable between the two scenarios, but there are several notable exceptions. First, the gap between 50 MeV and 5 GeV that is present for the hypercharge-coupled ALP is largely filled in the $\mathrm{SU}(2)_L$-coupled scenario because of powerful flavor constraints. There are a few gaps in the sensitivity of the *BABAR* analysis in the vicinity of the SM diphoton resonances (π^0, η, η'), but otherwise much of the low-mass parameter space is already ruled out. At high masses, dedicated LHC searches for ALPs decaying into photon pairs more powerfully constrain the hypercharge scenario than the $\mathrm{SU}(2)_L$ scenario; for the latter, the constraints at high mass predominantly come from WW searches, which suffer larger QCD backgrounds than the $\gamma\gamma$ searches.

We next consider the parameters motivated by the thermal DM abundance. In Figs. 5 and 6, we show the coupling needed to obtain the observed DM abundance for fixed DM masses as a function of the ALP mass. We truncate the DM relic curves when our treatment of the ALP as a propagating particle is no longer valid, $\Gamma_a/M_a > 0.5$.

In Fig. 5, we show results for the hypercharge-coupled scenario with $Q_Y = 1$ and 10, corresponding to "small" and "large" gauge couplings. For the "small" gauge coupling limit, which is largely unconstrained by indirect detection searches, we find that most of the DM-motivated parameter space is beyond current constraints. However, improvements in coupling constraints of one to two orders of magnitude at future experiments would allow either the discovery or exclusion of large portions of the electroweak axion portal parameter space. By contrast, for "large" gauge coupling, most of the parameter space is excluded by both indirect detection and accelerator searches except in the immediate vicinity of resonance. Even the resonantly enhanced parameter space could be tested with the next generation of experiments, although for much of this parameter space the ALP decay into DM is kinematically allowed and the ALP no longer has a substantial visible decay fraction. We use dashed lines to indicate the parts of the DM contours when $\mathrm{BF}(a \to \bar{\chi}\chi) > 50\%$.

In Fig. 6, we show results for the $\mathrm{SU}(2)_L$-coupled scenario. Once again, we have "small" and "large" gauge coupling limits of $T(R) = 1/2$ and 110, respectively. The DM relic curves are qualitatively similar to those for hypercharge coupling, except that much of the viable parameter space is ruled out for ALP masses below 5 GeV due to strong flavor constraints. This is true even

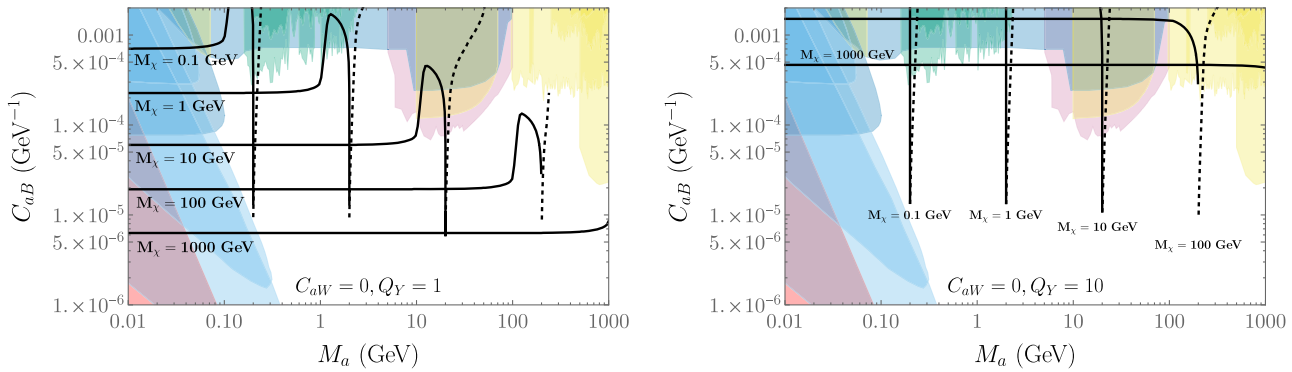

FIG. 5. Accelerator and collider bounds on the hypercharge-coupled scenario including thermal-relic contours for various DM masses. These contours assume (left) $Q_Y = 1$; (right) $Q_Y = 10$. The dashed line indicates the DM parameters for which $\mathrm{BF}(a \to \bar{\chi}\chi) > 50\%$. DM curves are truncated when our treatment of the ALP as a propagating particle is no longer valid, $\Gamma_a/M_a > 0.5$. There is also a sliver of parameter space excluded from supernovae constraints depicted in red.

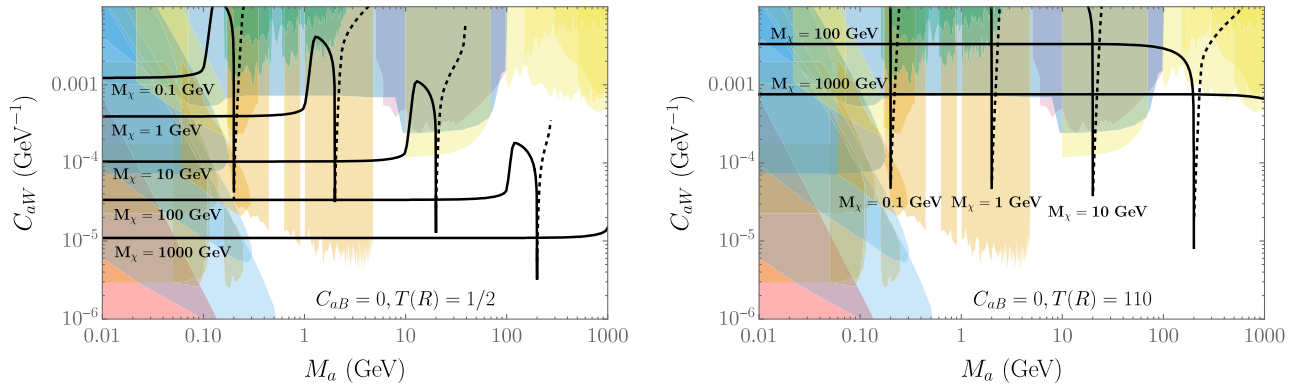

FIG. 6. Accelerator and collider bounds on the $SU(2)_L$-coupled scenario including thermal-relic contours for various DM masses. These contours assume (left) $T(R) = 1/2$; (right) $T(R) = 110$. The dashed line indicates the DM parameters for which $\mathrm{BF}(a \to \bar{\chi}\chi) > 50\%$. DM curves are truncated when our treatment of the ALP as a propagating particle is no longer valid, $\Gamma_a/M_a > 0.5$. There is also a sliver of parameter space excluded from supernovae constraints depicted in red.

in the immediate vicinity of resonances, provided the decay to visible states still dominates the ALP decay width.

Finally, we comment on the coupling peak (or, more accurately, dip) of the resonant annihilation region. As argued in Sec. III A, the coupling needed to obtain a thermal abundance is independent of M_a for annihilation slightly above resonance, but scales like $M_a^{-1/3}$ for annihilation exactly on resonance. Consequently, we find at lower ALP masses that the smallest allowed coupling near resonance is associated with annihilation just above resonance and the peaks are of similar height, modulated only by the fact that g_* changes steeply for $T \sim 0.1$–10 GeV and therefore the couplings predicted by freeze-out vary slightly due to this effect. At larger ALP masses, annihilation exactly on resonance becomes more important due to the $M_a^{-1/3}$ scaling of the coupling on resonance, and consequently deeper troughs are predicted at larger M_a.

VI. FREEZE-IN DARK MATTER

For sufficiently small couplings, DM never reaches thermal equilibrium and its abundance is established through freeze-in [95,96]. In this case, it is also possible that the ALP is out of equilibrium for some or most of the relevant cosmic history during which DM is produced. However, since ALPs are produced at $\mathcal{O}(g_{aVV}^2) \sim \frac{1}{f^2}$ and DM is produced at $\mathcal{O}(\frac{g_{aVV}^2}{f^2}) \sim \frac{1}{f^4}$, the ALP thermalizes in a large share of the DM freeze-in parameter space.

Because ALPs are pseudo-Goldstone bosons, their couplings arise from higher-dimensional operators. The additional degree of momentum-dependence in ALP-mediated reaction densities results in larger scattering rates relative to Hubble expansion at earlier times, and this leads to UV freeze-in [96,97]. Predictions of UV freeze-in models are sensitive to the reheat temperature, T_{RH}, although as we will see there are also freeze-in modes that occur predominantly in the infrared (IR).

We assume an instantaneous reheating such that the DM and ALP abundances are zero at T_{RH}. When freeze-in proceeds through an EFT operator of dimension-7 or lower, as is the case for the axion portal, the instantaneous approximation largely agrees with results from more careful implementations of reheating [98]. However, the approximation breaks down for masses close to or exceeding the reheat temperature, and so we only show numerical results for $M_\chi, M_a \leq T_{\mathrm{RH}}$.

Our findings qualitatively agree with the regimes identified in previous studies of ALP-portal freeze-in that made different coupling assumptions [20,21,62,63]. We restrict ourselves to the hypercharge-coupled scenario since we expect results with $SU(2)_L$ coupling to be similar. We find inconsistencies with the results from the freeze-in module of microMEGAs, however, and so we use our own Boltzmann equations to study freeze-in.

A. Calculation of DM abundance

We solve the same Boltzmann equations as before, although we relax the assumption that the ALP and/or DM begin in thermal equilibrium. This means we track the ALP abundance in addition to the DM abundance. DM can be directly produced in $\gamma\gamma \to \bar{\chi}\chi$ collisions via an intermediate virtual ALP and from within the hidden sector through ALP decays ($a \to \bar{\chi}\chi$) and scattering ($aa \to \bar{\chi}\chi$). ALPs are produced through inverse decays ($\gamma\gamma \to a$) and scattering processes (such as $e^+e^- \to \gamma a$ and $e^-\gamma \to e^- a$) involving SM gauge bosons. The latter are suppressed by α_{EM} relative to inverse decays, but their contribution to the ALP abundance is enhanced by T_{RH}/M_a, which can dominate for higher reheat temperatures.

The ALP production rate through $2 \to 2$ scattering processes depends on which charged species are present in the plasma and is therefore sensitive to various SM mass thresholds. A full treatment of all the mass-dependent effects of every species would be involved and would

likely have negligible impact on the results when the reheat temperature is far from any mass thresholds due to the UV dominance of scattering. We adopt a simplified treatment where we include only scattering involving charged SM fermions, both leptons and quarks ($\bar{f}f \to \gamma a$, $f\gamma \to fa$, and $\bar{f}\gamma \to \bar{f}a$). We neglect fermion masses in these rates except for when they act as regulators for t-channel logarithmic divergences, and we only include scattering involving a fermion f for temperatures $T \geq M_f$.[4] We have checked that using the massless rates for the reaction densities agrees well with the full rates for $T \gtrsim M_f$; the rates are in any case Boltzmann-suppressed for smaller temperatures. This treatment allows us to efficiently obtain results over a wide parameter space to accuracy in the couplings to better than $\mathcal{O}(1)$ precision.

Our Boltzmann equations assume that the only ALP-SM coupling is to photons: this is perfectly fine for reheat temperatures below the electroweak scale. For reheat temperatures well above the electroweak scale, the correct treatment in the hypercharge-coupled scenario replaces a photon with a hypercharge boson, and so we can simply replace $g_{a\gamma\gamma} \to C_{aB}$. In this limit, we additionally include $a \to \gamma Z$ and $a \to ZZ$ decays when calculating the ALP width.

We provide more details of the Boltzmann equations in Appendix A.

B. Regimes of Freeze-In

The dominant mode of DM production depends on the relative values of M_a, M_χ, and $T_{\rm RH}$, as well as on the relative strength of the ALP coupling to SM gauge bosons vs DM. We choose the same coupling benchmarks as outlined in Sec. II A. We adopt the following approach: we fix $T_{\rm RH}$ and M_χ and then study how the freeze-in coupling depends on M_a. We then repeat for different values of $T_{\rm RH}$ and M_χ.

Starting from $M_a \ll M_\chi$ and then gradually increasing M_a, we find the following regimes for freeze-in:

M_a below DM decay threshold: in this limit, $M_a < 2M_\chi$, the $a \to \bar{\chi}\chi$ decay is kinematically forbidden and the ALP always decays visibly. DM is produced via scattering at $\mathcal{O}(\frac{1}{f^4})$, and to obtain an appreciable DM abundance the coupling is large enough to thermalize a (whose production rate is $\sim\frac{1}{f^2}$). The dominant DM production mechanism is typically $aa \to \bar{\chi}\chi$ and occurs most rapidly just after reheating. If $T_{\rm RH} \gg M_a, M_\chi$, then this DM production cross section scales like $\frac{M_\chi^2}{f^4}$ and is independent of M_a. Consequently, we see an essentially flat dependence of the freeze-in coupling with increasing M_a in this regime, while $C_{aB} \sim \frac{1}{f}$ varies with DM mass like $M_\chi^{-1/2}$.

M_a just above DM decay threshold: as M_a increases above $2M_\chi$, the DM production channel $a \to \bar{\chi}\chi$ opens up. For our benchmarks, the ALP coupling to SM gauge bosons is smaller than that to DM, so the ALP does not need to be much above threshold before ${\rm BF}(a \to \bar{\chi}\chi)$ becomes nearly 100%. Since ALPs are produced at $\mathcal{O}(\frac{1}{f^2})$ and every ALP decays into a χ and a $\bar{\chi}$, the DM abundance is now also established at $\mathcal{O}(\frac{1}{f^2})$. Consequently, the only viable parameter space has both DM and the ALP out of equilibrium, and the coupling needed to achieve the correct DM abundance plummets by many orders of magnitude. The fact that the ALP decays invisibly in this regime also affects certain terrestrial and cosmological constraints on ALPs.

The predicted freeze-in coupling is largely independent of M_a in this limit because the ALP production rate is also largely independent of M_a, with a minor dependence on M_a arising from the fact that the inverse decay rate $\gamma\gamma \to a$ depends on M_a but is typically subdominant to $2 \to 2$ scattering (along with a mild logarithmic dependence on M_a for scattering processes with a t-channel singularity). Since $Y_\chi^{\rm after\ decay} = Y_a^{\rm before\ decay}$, the DM abundance remains fixed for fixed coupling independent of M_a.

The ALP is visible again: as M_a continues to increase, its visible branching fraction also increases until it dominates the decay width. To see why this is the case, note that ${\rm BF}(a \to \bar{\chi}\chi) \sim \frac{M_\chi^2}{M_a^2}$ and decreases with larger ALP masses. We now have $Y_\chi^{\rm after\ decay} = {\rm BF}(a \to \bar{\chi}\chi) Y_a^{\rm before\ decay}$, and consequently a larger ALP-photon coupling is needed for larger ALP masses to produce more ALPs in order to compensate for the diminished DM branching fraction. According to this argument, we predict a scaling $C_{aB} \sim M_a$. Somewhat counterintuitively, the ALP also decays visibly in experiments, even though its decay to DM is kinematically allowed. The value of M_a at which we enter this regime depends on the relative coupling strength to DM vs photons: a smaller coupling to photons implies larger values of M_a for the transition to visible decays.

The ALP reaches equilibrium: the scaling in the previous regime continues until the ALP reaches equilibrium at reheating. At this point, increasing the SM coupling to ALPs no longer increases the ALP abundance since Y_a saturates at $Y_a^{\rm eq}$. The result is a small range of ALP masses over which the freeze-in coupling can vary by many orders of magnitude since its effect on the ALP abundance (and consequently the DM abundance) has little effect in the approach to equilibrium.

[4]We use the kaon mass for the strange quark threshold and the pion mass for the up and down quark thresholds.

This transition region can appear as a sharp kink or a mild wiggle depending on the relative values of various parameters.

For $T_{\rm RH} \gg M_a$, the fact that the ALP reaches equilibrium at the time of reheating does not mean that it *stays* in equilibrium throughout its history since the $2 \to 2$ scattering rate falls sharply as the Universe cools. For sufficiently large couplings, it also maintains equilibrium at $T \sim M_a$ through (inverse) decays to photons. This qualitatively changes the DM prediction: whereas before we had $Y_\chi^{\text{after decay}} \propto Y_a^{\text{before decay}}$, now the ALP abundance is continually replenished from the SM bath as it decays into DM. The result is that the DM abundance depends on $\Gamma(a \to \bar{\chi}\chi)$ (*not* the branching fraction) integrated over time. We can approximate this integrated quantity as $\Gamma(a \to \bar{\chi}\chi)/H(T = M_a)$, which scales as $\frac{1}{f^2 M_a}$.

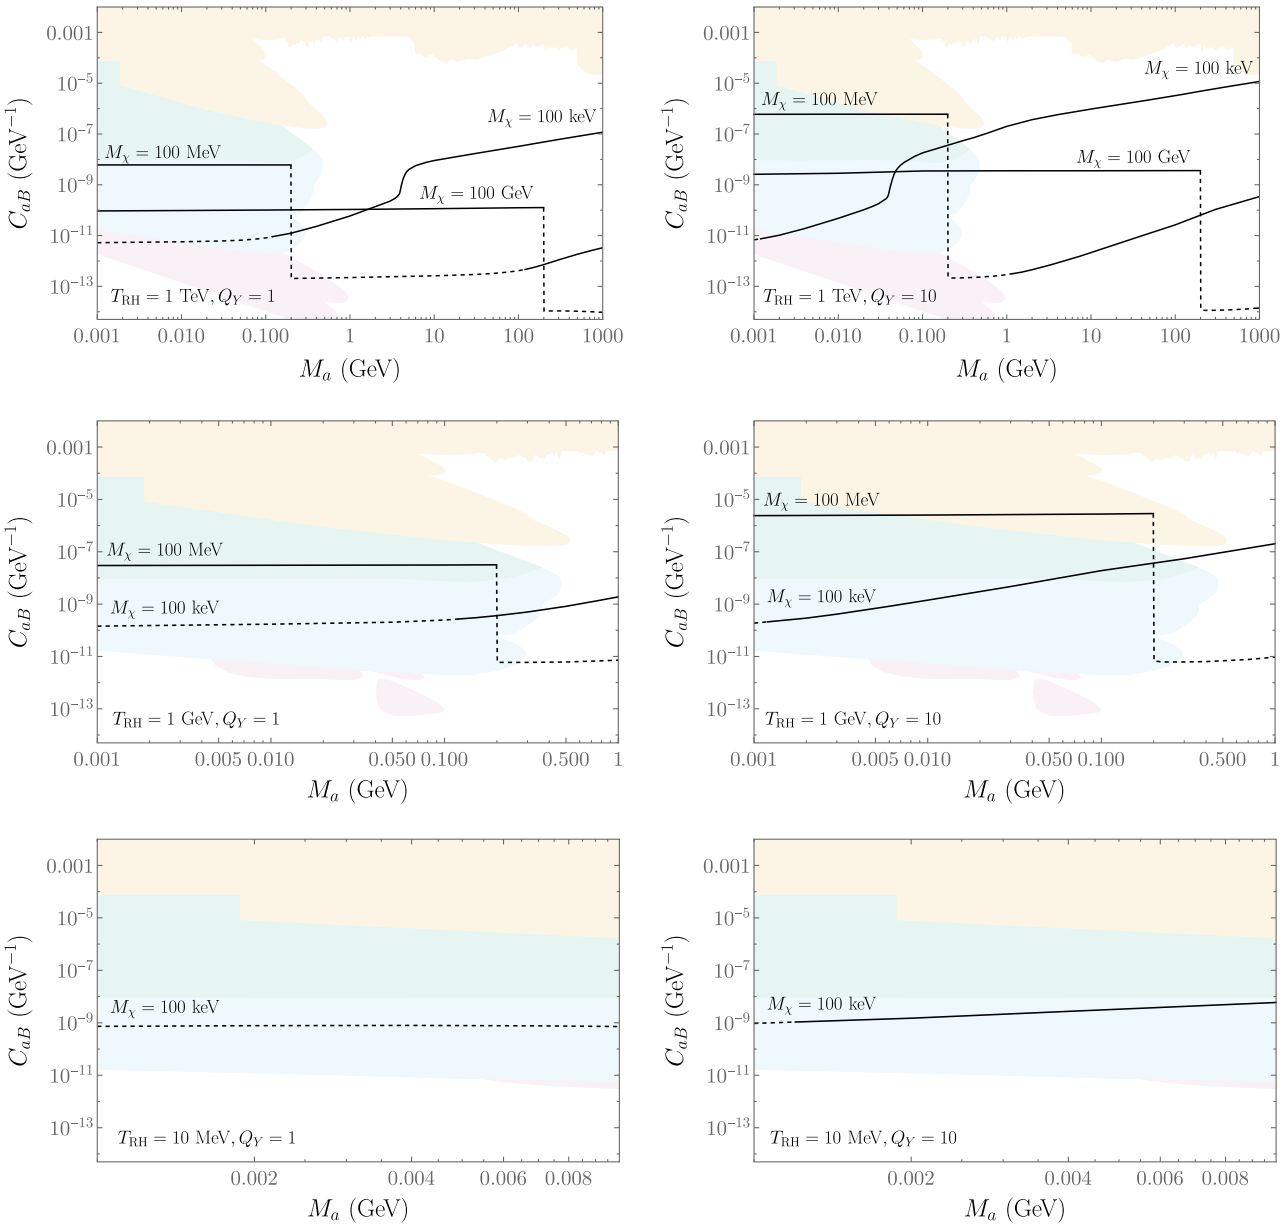

FIG. 7. Estimated couplings for DM freeze-in along with observational constraints. In all plots, $C_{aW} = 0$, and we take $Q_Y = 1$ ($Q_Y = 10$) in the left (right) column. We assume a different reheat temperature in each row: (top) 1 TeV; (middle) 1 GeV; (bottom) 10 MeV. In this figure, we consider only masses satisfying $M_\chi, M_a \leq T_{\rm RH}$. Cosmological constraints are shown in red, supernova cooling in green, gamma rays and axion-sourced fireballs from supernovae and neutron star mergers in blue, and accelerator/collider constraints in orange. The DM abundance curves are dashed when $\text{BF}(a \to \bar{\chi}\chi) > 50\%$; in these regions, most accelerator/collider, cosmological, and gamma-ray/fireball constraints are relaxed.

We therefore predict that the coupling $C_{aB} \sim \frac{1}{f} \sim \sqrt{M_a}$ in this regime.

The ALP mass exceeds the reheat temperature: when $M_a > T_{\rm RH}$, the DM production rate from ALP decays becomes Boltzmann suppressed. The result is that DM is produced from photon scattering via a virtual intermediate ALP with a maximum rate $\Gamma_{\gamma\gamma\to\bar{\chi}\chi} \sim \frac{C_{aB}^2 M_\chi^2 T_{\rm RH}^7}{f^2 M_a^4}$. Consequently, arbitrarily large values of C_{aB} are viable for freeze-in provided $T_{\rm RH} \ll M_a$. At the same time, the ALP EFT might break down in this limit and so care is needed to interpret any results. We do not quantitatively pursue this regime in any further detail but nevertheless point it out as a phenomenological possibility.

C. Numerical Results

To illustrate the regimes outlined in the previous section and to numerically compare the DM parameter space with current bounds, we select several reheat temperatures and DM masses, and for each combination of $T_{\rm RH}$ and M_χ we show the DM-motivated coupling as a function of the ALP mass. We consider reheat temperatures of 1 TeV, 1 GeV, and 10 MeV, with the latter approximating the lowest value of $T_{\rm RH}$ consistent with big bang nucleosynthesis (BBN). We study DM masses ranging from 100 keV (chosen to be consistent with structure-formation constraints on warm freeze-in DM [99]) to 100 GeV and ALP masses ranging from 1 MeV to 1 TeV. We present results for the hypercharge-coupled scenario with $Q_Y = 1$ and 10.

We show our results in Fig. 7. Collectively, each of the freeze-in regimes from the previous section is present in our results (for example: the sharp drop in coupling predicted at $M_a = 2M_\chi$, the sharp increase in coupling predicted well above $2M_\chi$ where the ALP equilibrates, the various coupling-mass scaling relations, etc.). We see, as expected, that the range of ALP masses for which $\mathrm{BF}(a \to \bar{\chi}\chi) \approx 100\%$ is much smaller for $Q_Y = 10$ than for $Q_Y = 1$, and consequently the value of M_a at which the ALP comes into equilibrium above the DM threshold is lower at higher Q_Y. It is evident that freeze-in DM is consistent with a very wide range of DM masses, ALP masses, and coupling strengths.

We also show constraints on the ALP assuming that it decays visibly, which is true for most ALP masses even above the DM threshold. The accelerator and collider constraints are the same as those derived in Sec. V. We impose constraints from BBN on primordial element abundances [100] and from Planck on new relativistic degrees of freedom [78]; these were derived in Ref. [101] for ALP freeze-in with various reheat temperatures. We also include astrophysical constraints from supernova cooling [91], gamma rays, and fireballs from ALP bursts produced in supernovae [102–105], low-energy supernova [106], and axion-sourced fireballs produced during neutron stars mergers [107]. We use the tabulated bounds in AXIONLIMITS when applying astrophysical constraints [108]. Note that most of the constraints apart from supernova cooling are relaxed when the ALP decays invisibly. Indirect-detection bounds are not typically strong enough to constrain the freeze-in parameter space.

Taken together, we see that the freeze-in scenario is somewhat constrained for ALP masses below 100 MeV, although there is still viable parameter space depending on the DM mass and reheat temperature, especially when the ALP decays invisibly. For larger ALP masses, there are essentially no constraints on freeze-in: the predicted couplings are sufficiently small that it will be challenging for even future experiments to achieve sensitivity.

VII. DISCUSSION AND CONCLUSIONS

We have systematically studied the cosmology and phenomenology of electroweak ALP portal DM. For thermal DM, we find the parameter space already meaningfully constrained by a combination of indirect detection of DM annihilation into photon lines and direct experimental searches for ALPs. There remains a substantial viable parameter space, much of which will be tested in future experiments. In particular, closing the well-known MeV-GeV gap in gamma ray searches could significantly improve sensitivity to DM masses in this range.

Interestingly, over much of the thermal parameter space we find that the validity of the EFT is somewhat marginal, with the PQ-breaking scale typically constrained to $f \lesssim 1$ TeV. The radial mode(s) of the PQ-breaking sector and heavy fermions mediating ALP-gauge interactions, all of which we have neglected, could play an appreciable role in the model's cosmology and phenomenology. This is particularly important for LHC searches since deviations from the EFT will first appear at high-energy colliders. A future detailed study is warranted of how ALP phenomenology changes in UV completions of the model, along with the interplay between ALP constraints and direct searches for new heavy states.

We have also investigated the possibility of freeze-in DM, finding that some parameter space, particularly at low ALP masses, is constrained by a combination of astrophysical and cosmological bounds. However, the predicted couplings are sufficiently small that there are few meaningful constraints on freeze-in for ALP masses above 100 MeV. We adopted a simple, instantaneous model of reheating, and it is possible that some of our DM-motivated parameters would change with a more realistic treatment of reheating.

ACKNOWLEDGMENTS

We are grateful to Benjamin Khoury and Chris Ranlett for collaboration at early stages of this work. We thank Eder

Izaguirre, Tongyan Lin, and Dave Tucker-Smith for helpful conversations. We are grateful to Edoardo Vitagliano for his helpful comments on astrophysical constraints. This work is supported by Research Corporation for Science Advancement through Cottrell Scholar Grant No. 27632 and by the Campbell Fund of Harvey Mudd College. B .S. is grateful for the hospitality of Perimeter Institute, where part of this work was carried out. Research at Perimeter Institute is supported in part by the Government of Canada through the Department of Innovation, Science and Economic Development and by the Province of Ontario through the Ministry of Colleges and Universities.

APPENDIX A: BOLTZMANN EQUATIONS

In this Appendix, we present the coupled Boltzmann equations for our model. The cross section and decay rates can be found in Appendix B.

1. Freeze-out

For thermal DM, most of the DM-coupling curve is derived using microMEGAs. However, we use the Boltzmann equations presented below to validate results from microMEGAs. We also use solutions to our Boltzmann equations in the vicinity of the ALP resonance due to greater numerical stability. We include only ALP-photon couplings, although in Sec. III A we have described how the cross section on resonance can be modified to take into account decays to weak gauge bosons as well.

For thermal DM, we can assume the ALP is always in equilibrium, while χ is in kinetic (but not necessarily chemical) equilibrium. We use Maxwell-Boltzmann statistics and neglect quantum statistical factors in the Boltzmann equation. We use $Y_\chi \equiv n_\chi/s$ to track the χ comoving number density and $x \equiv M_\chi/T$ as a dimensionless time variable. The Boltzmann equation for χ is:

$$\begin{aligned}\frac{dY_\chi}{dx} = &-\frac{s\langle\sigma_{\bar{\chi}\chi\to\gamma\gamma}v\rangle^{\rm sub}(Y_\chi^{\rm eq})^2}{xH(x)}\left[\frac{Y_\chi^2}{(Y_\chi^{\rm eq})^2}-1\right]\\ &-\frac{s\langle\sigma_{\bar{\chi}\chi\to aa}v\rangle(Y_\chi^{\rm eq})^2}{xH(x)}\left[\frac{Y_\chi^2}{(Y_\chi^{\rm eq})^2}-1\right]\\ &-\frac{\langle\Gamma_a\rangle Y_a^{\rm eq}}{xH(x)}{\rm BF}(a\to\bar{\chi}\chi)\left[\frac{Y_\chi^2}{(Y_\chi^{\rm eq})^2}-1\right]. \qquad \text{(A1)}\end{aligned}$$

In deriving these equations, we have used the real-intermediate-state (RIS) subtraction scheme [109–111] to remove double-counting from the portion of the $\bar{\chi}\chi\to\gamma\gamma$ process in which the ALP is on shell,

$$\begin{aligned}&\langle\sigma_{\bar{\chi}\chi\to\gamma\gamma}v\rangle^{\rm sub}\\ &=\langle\sigma_{\bar{\chi}\chi\to\gamma\gamma}v\rangle-\frac{Y_a^{\rm eq}}{s(Y_\chi^{\rm eq})^2}\langle\Gamma_a\rangle{\rm BF}(a\to\bar{\chi}\chi){\rm BF}(a\to\gamma\gamma). \qquad \text{(A2)}\end{aligned}$$

2. Freeze-in

For freeze-in, we can no longer assume the ALP is in equilibrium. We now solve a system of Boltzmann equations for $Y_a(x)$ and $Y_\chi(x)$. Because $2\to2$ scattering predominantly occurs in the UV at $T\sim T_{\rm RH}$ while inverse decays into ALPs occur in the IR at $T\sim M_a$, the phase space of these processes does not appreciably overlap and this eliminates the need for the RIS subtraction. The Boltzmann equations are:

$$\begin{aligned}\frac{dY_a}{dx} = &-\frac{\langle\Gamma_a\rangle Y_a^{\rm eq}}{xH(x)}{\rm BF}(a\to\gamma\gamma)\left(\frac{Y_a}{Y_a^{\rm eq}}-1\right)\\ &-\frac{\langle\Gamma_a\rangle Y_a^{\rm eq}}{xH(x)}{\rm BF}(a\to\bar{\chi}\chi)\left[\frac{Y_a}{Y_a^{\rm eq}}-\frac{Y_\chi^2}{(Y_\chi^{\rm eq})^2}\right] \qquad \text{(A3)}\end{aligned}$$

$$\begin{aligned}&+\frac{2s\langle\sigma_{\bar{\chi}\chi\to aa}v\rangle(Y_\chi^{\rm eq})^2}{xH(x)}\left[\frac{Y_\chi^2}{(Y_\chi^{\rm eq})^2}-\frac{Y_a^2}{(Y_a^{\rm eq})^2}\right]\\ &-\frac{s\langle\sigma_{{\rm SM}a\to{\rm SM\,SM}}v\rangle Y_{\rm SM}^{\rm eq}Y_a^{\rm eq}}{xH(x)}\left(\frac{Y_a}{Y_a^{\rm eq}}-1\right), \qquad \text{(A4)}\end{aligned}$$

$$\begin{aligned}\frac{dY_\chi}{dx} = &-\frac{s\langle\sigma_{\bar{\chi}\chi\to\gamma\gamma}v\rangle(Y_\chi^{\rm eq})^2}{xH(x)}\left[\frac{Y_\chi^2}{(Y_\chi^{\rm eq})^2}-1\right]\\ &-\frac{s\langle\sigma_{\bar{\chi}\chi\to aa}v\rangle(Y_\chi^{\rm eq})^2}{xH(x)}\left[\frac{Y_\chi^2}{(Y_\chi^{\rm eq})^2}-\frac{Y_a^2}{(Y_a^{\rm eq})^2}\right] \qquad \text{(A5)}\end{aligned}$$

$$-\frac{\langle\Gamma_a\rangle Y_a^{\rm eq}}{xH(x)}{\rm BF}(a\to\bar{\chi}\chi)\left[\frac{Y_\chi^2}{(Y_\chi^{\rm eq})^2}-\frac{Y_a}{Y_a^{\rm eq}}\right]. \qquad \text{(A6)}$$

For the production of ALPs from SM $2\to2$ scattering, we include the processes $f\bar{f}\to\gamma a$, $f\gamma\to fa$, and $\bar{f}\gamma\to\bar{f}a$ for all charged SM fermions. These rates are weighted by appropriate squared-charge and color factors, and at temperature T we include all charged fermion species with masses below T.

For $T_{\rm RH}=10$ MeV and 1 GeV, we can neglect interactions with the weak gauge bosons. For $T_{\rm RH}=1$ TeV, however, we must include these effects. Since we only consider the hypercharge-coupled scenario for freeze-in, and the $2\to2$ scattering rates above are UV dominated, we can simply replace the photon with a hypercharge boson in the above rates, which is equivalent to taking $g_{a\gamma\gamma}\to C_{aB}$. For (inverse) ALP decays, we include decay rates to all SM boson pairs.

APPENDIX B: CROSS SECTIONS AND DECAY RATES

Here, we present various decay rates and cross sections implemented in our Boltzmann equations.

The ALP partial widths are

$$\Gamma(a \to \bar{\chi}\chi) = \frac{M_a M_\chi^2}{8\pi f^2}\sqrt{1-\frac{4M_\chi^2}{M_a^2}}, \tag{B1}$$

$$\Gamma(a \to \gamma\gamma) = \frac{g_{a\gamma\gamma}^2 M_a^3}{64\pi}, \tag{B2}$$

$$\Gamma(a \to \gamma Z) = \frac{g_{a\gamma Z}^2 M_a^3}{32\pi}\left(1-\frac{M_Z^2}{M_a^2}\right)^3, \tag{B3}$$

$$\Gamma(a \to ZZ) = \frac{g_{aZZ}^2 M_a^3}{64\pi}\left(1-\frac{4M_Z^2}{M_a^2}\right)^{3/2}, \tag{B4}$$

$$\Gamma(a \to W^+W^-) = \frac{g_{aWW}^2 M_a^3}{32\pi}\left(1-\frac{4M_W^2}{M_a^2}\right)^{3/2} \tag{B5}$$

for the decay modes shown in Fig. 8. All couplings are defined in Sec. II. The thermally averaged decay width is related to the zero-temperature decay width via

$$\langle\Gamma\rangle = \Gamma\frac{K_1(M_a/T)}{K_2(M_a/T)}, \tag{B6}$$

where $K_1(x)$ and $K_2(x)$ are modified Bessel functions of the second kind.

DM can annihilate into a pair of SM bosons via an on- or off-shell axion as shown in Fig. 9. Its cross section in the narrow-width approximation is

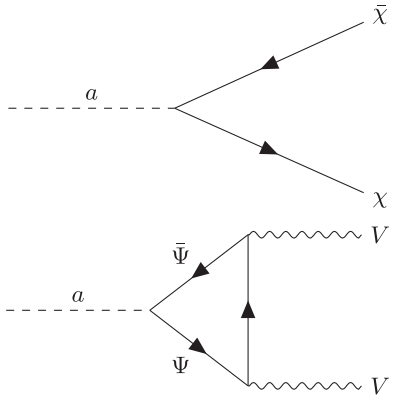

FIG. 8. *Top:* ALP decay to a pair of dark matter and anti-dark matter. *Bottom:* ALP decay into a pair of electroweak bosons $VV \in \{\gamma\gamma, W^+W^-, ZZ, \gamma Z\}$ via a loop of massive charged fermions Ψ. In the effective field theory, the loop can be collapsed into the aVV vertex with coupling strength g_{aVV}.

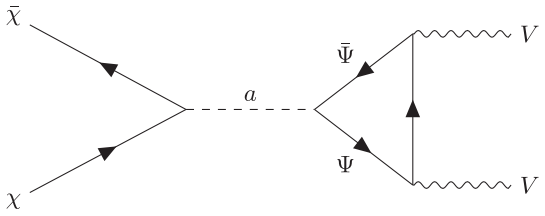

FIG. 9. DM annihilates into a pair of electroweak bosons $VV \in \{\gamma\gamma, W^+W^-, ZZ, \gamma Z\}$ via an on or off shell axion.

$$\sigma_{\bar{\chi}\chi\to\gamma\gamma} = \frac{g_{a\gamma\gamma}^2 M_\chi^2}{128\pi f^2}\frac{1}{\sqrt{1-4M_\chi^2/s}}\frac{s^2}{(s-M_a^2)^2+M_a^2\Gamma_a^2}, \tag{B7}$$

where s here is the Mandelstam variable, not the entropy density.

The cross section for DM annihilation to ALPs is

$$\sigma_{\bar{\chi}\chi\to aa} = \frac{M_\chi^2}{64\pi f^4 s^2}\frac{1}{1-4M_\chi^2/s}\left[\frac{M_\chi^2 s^2 - 4M_a^2M_\chi^2 s + M_a^4(s-2M_\chi^2)}{M_\chi^2 s - 4M_a^2M_\chi^2 + M_a^4}\sqrt{s-4M_a^2}\sqrt{s-4M_\chi^2}\right.$$
$$\left. - \frac{2M_\chi^2(s^2-4M_a^2s^2+2M_a^4)}{s-2M_a^2}\log\left(\frac{s-2M_a^2+\sqrt{s-4M_a^2}\sqrt{s-4M_\chi^2}}{s-2M_a^2-\sqrt{s-4M_a^2}\sqrt{s-4M_\chi^2}}\right)\right], \tag{B8}$$

with the relevant Feynman diagrams illustrated in Fig. 10. The squared matrix elements for some of the processes were calculated with the assistance of FeynCalc9.3 [112].

In both cases, the thermally averaged cross section is found from the cross section using [113]

$$\langle\sigma_{\bar{\chi}\chi}v\rangle = \frac{1}{8M_\chi^4 T K_2(M_\chi/T)^2}\int_{4M_\chi^2}^{\infty}\sigma_{\bar{\chi}\chi}\sqrt{s}(s-4M_\chi^2)K_1(\sqrt{s}/T)ds. \tag{B9}$$

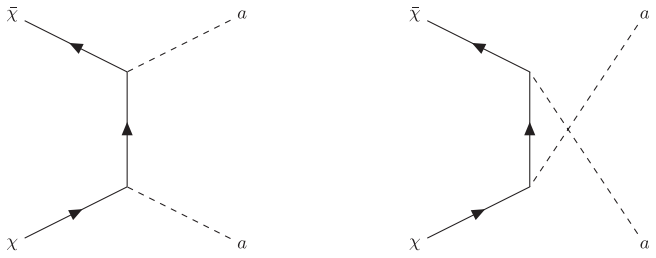

FIG. 10. Two Feynman diagrams contributing to DM annihilating into a pair of ALPs.

For freeze-in, we also need to calculate the rates of ALP production. In the limit of approximately massless ALP and fermions, we have

$$\sigma_{\gamma a\to e^+e^-} = \frac{\alpha g_{a\gamma\gamma}^2}{12}, \tag{B10}$$

$$\sigma_{e^-a\to e^-\gamma} = \frac{\alpha g_{a\gamma\gamma}^2}{2}\log\left(\frac{s^{3/2}}{M_e M_a^2}\right). \tag{B11}$$

The thermally averaged cross sections for these processes in the massless limit are found using:

$$\langle\sigma_{\mathrm{SM}a}v\rangle = \frac{1}{32T^5}\int_0^\infty \sigma_{\mathrm{SM}a}s^{3/2}K_1(\sqrt{s}/T)ds. \tag{B12}$$

The thermal average integrals can be evaluated in closed form, giving

$$\langle\sigma_{\gamma a\to e^+e^-}v\rangle = \frac{\alpha g_{a\gamma\gamma}^2}{12}, \tag{B13}$$

$$\langle\sigma_{e^-a\to e^-\gamma}v\rangle \approx \frac{\alpha g_{a\gamma\gamma}^2}{2}\log\left(\frac{60.2T^3}{M_e M_a^2}\right). \tag{B14}$$

We include SM fermion thermal masses in the log-divergent term of the cross section when they dominate over tree-level masses:

$$\frac{M_{Q_\mathrm{L}}^2}{T^2} = \frac{g_3^2}{3} + \frac{3g_2^2}{16} + \frac{g_1^2}{144}, \tag{B15}$$

$$\frac{M_{u_\mathrm{R}}^2}{T^2} = \frac{g_3^2}{3} + \frac{g_1^2}{9}, \tag{B16}$$

$$\frac{M_{d_\mathrm{R}}^2}{T^2} = \frac{g_3^2}{3} + \frac{g_1^2}{36}, \tag{B17}$$

$$\frac{M_{L_\mathrm{L}}^2}{T^2} = \frac{3g_2^2}{16} + \frac{g_1^2}{16}, \tag{B18}$$

$$\frac{M_{e_\mathrm{R}}^2}{T^2} = \frac{g_1^2}{4}, \tag{B19}$$

where g_3, g_2, and g_1 are the $\mathrm{SU}(3)_c$, $\mathrm{SU}(2)_L$, and $U(1)_Y$ gauge couplings, respectively. We neglect Yukawa couplings in these calculations: they are only relevant for the top quark, which gives only a small contribution to the overall ALP production rate.

APPENDIX C: A UV COMPLETION OF ALP-DM INTERACTION

We consider a massless Dirac fermion DM field, χ and a complex scalar field, Φ. The theory is classically invariant under a global axial $\mathrm{U}(1)_{\mathrm{PQ}}$ transformation given by $\chi \to e^{-i\theta\gamma^5}\chi$ and $\Phi \to e^{-2i\theta}\Phi$, with interaction

$$\mathcal{L} = -y\Phi\bar{\chi}_\mathrm{R}\chi_\mathrm{L} + \mathrm{H.c.} \tag{C1}$$

The symmetry is spontaneously broken with $\langle\Phi\rangle = \frac{f}{\sqrt{2}}$, and this gives DM a mass, $M_\chi = \frac{yf}{\sqrt{2}}$. If the symmetry is exact, then there exists a massless Goldstone boson, a, along with a massive radial mode, φ. It is convenient to parametrize the complex scalar field as $\Phi = \frac{(f+\varphi)}{\sqrt{2}}e^{\frac{ia}{f}}$, such that the shift symmetry $a \to a - 2\theta f$ is manifestly preserved and so that both real scalar fields have canonical normalization. If the symmetry is explicitly broken, then a is not exactly a Goldstone boson, but this decomposition remains approximately aligned with the mass eigenstates provided $M_a \ll f$.

When $M_\chi \ll M_\varphi$, there exists an EFT that includes χ and a with Lagrangian

$$\mathcal{L} = -M_\chi e^{\frac{ia}{f}}\bar{\chi}_\mathrm{R}\chi_\mathrm{L} + \mathrm{H.c.} \tag{C2}$$

The $e^{\frac{ia}{f}}$ can be removed by a field redefinition, $\chi \to e^{\frac{ia}{2f}\gamma^5}\chi$, in which case the ALP is removed from the Yukawa coupling in Eq. (C2) but then reappears in the DM kinetic term. After expanding out derivatives, we obtain the standard ALP-DM derivative coupling

$$\mathcal{L} = -\frac{1}{2f}\partial_\mu a\bar{\chi}\gamma^\mu\gamma^5\chi. \tag{C3}$$

This manifestly preserves the ALP shift symmetry with a single ALP-DM coupling.

Alternatively, we could have worked directly with Eq. (C2) by expanding $e^{\frac{ia}{f}}$ as a Taylor series. The leading term in the result is $i\frac{M_\chi}{f}a\bar{\chi}\gamma^5\chi$, although we note that this term on its own does not obey the shift symmetry. Indeed, we need an infinite sum to fully preserve the shift symmetry, although in practice we can truncate the sum at whatever order in the EFT that we are working. For example, if we wish to calculate amplitudes to $\mathcal{O}(\frac{1}{f^2})$ using

nonderivative couplings, then we must keep powers of a in the exponential up to quadratic, which results in an additional contact interaction $\frac{M_\chi}{2f^2}a^2\bar{\chi}\chi$.

The results obtained using either method are identical, but we find it simpler to work with the single derivative interaction in Eq. (C3).

There exist additional terms obtained by integrating out the radial mode, φ. For example, there exists an $a-a-\varphi$ vertex, which induces an operator $\frac{y}{\sqrt{2}fm_\varphi^2}\partial_\mu a\partial^\mu a\bar{\chi}\chi$ in the EFT, although this is a higher-dimension operator than the leading terms giving rise to $\bar{\chi}\chi\to aa$ annihilation and can therefore be neglected in our analysis.

[1] R. D. Peccei and H. R. Quinn, Phys. Rev. Lett. **38**, 1440 (1977).
[2] R. D. Peccei and H. R. Quinn, Phys. Rev. D **16**, 1791 (1977).
[3] S. Weinberg, Phys. Rev. Lett. **40**, 223 (1978).
[4] F. Wilczek, Phys. Rev. Lett. **40**, 279 (1978).
[5] E. Witten, Phys. Lett. **149B**, 351 (1984).
[6] J. P. Conlon, J. High Energy Phys. 05 (2006) 078.
[7] P. Svrcek and E. Witten, J. High Energy Phys. 06 (2006) 051.
[8] A. Arvanitaki, S. Dimopoulos, S. Dubovsky, N. Kaloper, and J. March-Russell, Phys. Rev. D **81**, 123530 (2010).
[9] J. M. Frere, D. R. T. Jones, and S. Raby, Nucl. Phys. **B222**, 11 (1983).
[10] A. E. Nelson and N. Seiberg, Nucl. Phys. **B416**, 46 (1994).
[11] J. Bagger, E. Poppitz, and L. Randall, Nucl. Phys. **B426**, 3 (1994).
[12] P. W. Graham, D. E. Kaplan, and S. Rajendran, Phys. Rev. Lett. **115**, 221801 (2015).
[13] Y. Nomura and J. Thaler, Phys. Rev. D **79**, 075008 (2009).
[14] M. Freytsis and Z. Ligeti, Phys. Rev. D **83**, 115009 (2011).
[15] M. J. Dolan, F. Kahlhoefer, C. McCabe, and K. Schmidt-Hoberg, J. High Energy Phys. 03 (2015) 171; 07 (2015) 103(E).
[16] K. Kaneta, H.-S. Lee, and S. Yun, Phys. Rev. D **95**, 115032 (2017).
[17] A. Kamada, H. Kim, and T. Sekiguchi, Phys. Rev. D **96**, 016007 (2017).
[18] Y. Hochberg, E. Kuflik, R. Mcgehee, H. Murayama, and K. Schutz, Phys. Rev. D **98**, 115031 (2018).
[19] S. Gola, S. Mandal, and N. Sinha, Int. J. Mod. Phys. A **37**, 2250131 (2022).
[20] P. J. Fitzpatrick, Y. Hochberg, E. Kuflik, R. Ovadia, and Y. Soreq, Phys. Rev. D **108**, 075003 (2023).
[21] J. A. Dror, S. Gori, and P. Munbodh, J. High Energy Phys. 09 (2023) 128.
[22] G. Armando, P. Panci, J. Weiss, and R. Ziegler, Phys. Rev. D **109**, 055029 (2024).
[23] M. J. Dolan, T. Ferber, C. Hearty, F. Kahlhoefer, and K. Schmidt-Hoberg, J. High Energy Phys. 12 (2017) 094; 03 (2021) 190(E).
[24] J. D. Bjorken, S. Ecklund, W. R. Nelson, A. Abashian, C. Church, B. Lu, L. W. Mo, T. A. Nunamaker, and P. Rassmann, Phys. Rev. D **38**, 3375 (1988).
[25] J. Blumlein *et al.*, Z. Phys. C **51**, 341 (1991).
[26] M. Acciarri *et al.* (L3 Collaboration), Phys. Lett. B **345**, 609 (1995).
[27] M. Acciarri *et al.* (L3 Collaboration), Phys. Lett. B **353**, 136 (1995).
[28] M. Acciarri *et al.* (L3 Collaboration), Phys. Lett. B **412**, 201 (1997).
[29] A. V. Artamonov *et al.* (E949 Collaboration), Phys. Lett. B **623**, 192 (2005).
[30] J. Abdallah *et al.* (DELPHI Collaboration), Eur. Phys. J. C **60**, 17 (2009).
[31] E. Abouzaid *et al.* (KTeV Collaboration), Phys. Rev. D **77**, 112004 (2008).
[32] I. Larin *et al.* (PrimEx Collaboration), Phys. Rev. Lett. **106**, 162303 (2011).
[33] S. Chatrchyan *et al.* (CMS Collaboration), Phys. Rev. Lett. **108**, 111801 (2012).
[34] G. Aad *et al.* (ATLAS Collaboration), Phys. Lett. B **710**, 538 (2012).
[35] G. Aad *et al.* (ATLAS Collaboration), J. High Energy Phys. 01 (2013) 086.
[36] J. Jaeckel, M. Jankowiak, and M. Spannowsky, Phys. Dark Universe **2**, 111 (2013).
[37] T. A. Aaltonen *et al.* (CDF Collaboration), Phys. Rev. Lett. **112**, 111803 (2014).
[38] C. Lazzeroni *et al.* (NA62 Collaboration), Phys. Lett. B **732**, 65 (2014).
[39] G. Aad *et al.* (ATLAS Collaboration), Eur. Phys. J. C **76**, 210 (2016).
[40] J. Jaeckel and M. Spannowsky, Phys. Lett. B **753**, 482 (2016).
[41] S. Knapen, T. Lin, H. K. Lou, and T. Melia, Phys. Rev. Lett. **118**, 171801 (2017).
[42] E. Izaguirre, T. Lin, and B. Shuve, Phys. Rev. Lett. **118**, 111802 (2017).
[43] M. Bauer, M. Neubert, and A. Thamm, J. High Energy Phys. 12 (2017) 044.
[44] M. Aaboud *et al.* (ATLAS Collaboration), J. High Energy Phys. 03 (2018) 042.
[45] A. M. Sirunyan *et al.* (CMS Collaboration), Phys. Lett. B **797**, 134826 (2019).
[46] J. K. Ahn *et al.* (KOTO Collaboration), Phys. Rev. Lett. **122**, 021802 (2019).
[47] D. Aloni, C. Fanelli, Y. Soreq, and M. Williams, Phys. Rev. Lett. **123**, 071801 (2019).
[48] G. Aad *et al.* (ATLAS Collaboration), J. High Energy Phys. 03 (2021) 243; 11 (2021) 050(E).

[49] F. Abudinén *et al.* (Belle-II Collaboration), Phys. Rev. Lett. **125**, 161806 (2020).
[50] S. Gori, G. Perez, and K. Tobioka, J. High Energy Phys. 08 (2020) 110.
[51] D. Banerjee *et al.* (NA64 Collaboration), Phys. Rev. Lett. **125**, 081801 (2020).
[52] E. Cortina Gil *et al.* (NA62 Collaboration), J. High Energy Phys. 03 (2021) 058.
[53] J. P. Lees *et al.* (*BABAR* Collaboration), Phys. Rev. Lett. **128**, 131802 (2022).
[54] A. Tumasyan *et al.* (CMS Collaboration), Phys. Rev. Lett. **127**, 261804 (2021).
[55] A. Tumasyan *et al.* (CMS Collaboration), Phys. Rev. D **105**, 032008 (2022).
[56] G. Aad *et al.* (ATLAS Collaboration), Phys. Lett. B **822**, 136651 (2021).
[57] M. Ablikim *et al.* (BESIII Collaboration), Phys. Lett. B **838**, 137698 (2023).
[58] J. Jerhot, B. Döbrich, F. Ertas, F. Kahlhoefer, and T. Spadaro, J. High Energy Phys. 07 (2022) 094.
[59] A. Mitridate, M. Papucci, C. W. Wang, C. Peña, and S. Xie, Phys. Rev. D **108**, 055040 (2023).
[60] G. Aad *et al.* (ATLAS Collaboration), J. High Energy Phys. 07 (2023) 234.
[61] A. Tumasyan *et al.* (CMS, TOTEM Collaborations), Phys. Rev. D **110**, 012010 (2024).
[62] A. Bharucha, F. Brümmer, N. Desai, and S. Mutzel, J. High Energy Phys. 02 (2023) 141.
[63] D. K. Ghosh, A. Ghoshal, and S. Jeesun, J. High Energy Phys. 01 (2024) 026.
[64] M. Bauer, M. Neubert, S. Renner, M. Schnubel, and A. Thamm, J. High Energy Phys. 04 (2021) 063.
[65] M. Bauer, M. Neubert, S. Renner, M. Schnubel, and A. Thamm, J. High Energy Phys. 09 (2022) 056.
[66] A. Belyaev, N. D. Christensen, and A. Pukhov, Comput. Phys. Commun. **184**, 1729 (2013).
[67] A. Alloul, N. D. Christensen, C. Degrande, C. Duhr, and B. Fuks, Comput. Phys. Commun. **185**, 2250 (2014).
[68] G. Bélanger, F. Boudjema, A. Goudelis, A. Pukhov, and B. Zaldivar, Comput. Phys. Commun. **231**, 173 (2018).
[69] A. Albert, G. A. Gomez-Vargas, M. Grefe, C. Munoz, C. Weniger, E. D. Bloom, E. Charles, M. N. Mazziotta, and A. Morselli (Fermi-LAT Collaboration), J. Cosmol. Astropart. Phys. 10 (2014) 023.
[70] J. W. Foster, Y. Park, B. R. Safdi, Y. Soreq, and W. L. Xu, Phys. Rev. D **107**, 103047 (2023).
[71] H. Abe *et al.* (MAGIC Collaboration), Phys. Rev. Lett. **130**, 061002 (2023).
[72] H. Abdallah *et al.* (HESS Collaboration), Phys. Rev. Lett. **120**, 201101 (2018).
[73] A. R. Pullen, R.-R. Chary, and M. Kamionkowski, Phys. Rev. D **76**, 063006 (2007); **83**, 029904(E) (2011).
[74] K. K. Boddy and J. Kumar, Phys. Rev. D **92**, 023533 (2015).
[75] A. De Angelis *et al.* (e-ASTROGAM Collaboration), Exp. Astron. **44**, 25 (2017).
[76] J. F. Navarro, C. S. Frenk, and S. D. M. White, Astrophys. J. **462**, 563 (1996).
[77] M. Cirelli, G. Corcella, A. Hektor, G. Hutsi, M. Kadastik, P. Panci, M. Raidal, F. Sala, and A. Strumia, J. Cosmol. Astropart. Phys. 03 (2011) 051; 10 (2012) E01.
[78] N. Aghanim *et al.* (Planck Collaboration), Astron. Astrophys. **641**, A6 (2020); **652**, C4(E) (2021).
[79] T. R. Slatyer, Phys. Rev. D **93**, 023527 (2016).
[80] D. Egana-Ugrinovic, M. Low, and J. T. Ruderman, J. High Energy Phys. 05 (2018) 012.
[81] M. M. Altakach, P. Lamba, R. Masełek, V. A. Mitsou, and K. Sakurai, Eur. Phys. J. C **82**, 848 (2022).
[82] A. Hayrapetyan *et al.* (CMS Collaboration), arXiv:2402.09932.
[83] M. T. Frandsen, U. Haisch, F. Kahlhoefer, P. Mertsch, and K. Schmidt-Hoberg, J. Cosmol. Astropart. Phys. 10 (2012) 033.
[84] L. Darmé, F. Giacchino, E. Nardi, and M. Raggi, J. High Energy Phys. 06 (2021) 009.
[85] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, J. High Energy Phys. 07 (2014) 079.
[86] L. Darmé *et al.*, Eur. Phys. J. C **83**, 631 (2023).
[87] M. Aaboud *et al.* (ATLAS Collaboration), J. High Energy Phys. 10 (2017) 112.
[88] M. Ablikim *et al.* (BESIII Collaboration), Phys. Rev. D **110**, L031101 (2024).
[89] S. Adler *et al.* (E787 Collaboration), Phys. Rev. D **70**, 037102 (2004).
[90] A. V. Artamonov *et al.* (BNL-E949 Collaboration), Phys. Rev. D **79**, 092004 (2009).
[91] F. Ertas and F. Kahlhoefer, J. High Energy Phys. 07 (2020) 050.
[92] Search for axion-like particles in photonic final states with the FASER detector at the LHC, Technical Report No. CERN-FASER-CONF-2024-001, CERN, Geneva, 2024.
[93] R. Aaij *et al.* (LHCb Collaboration), Phys. Rev. Lett. **115**, 161802 (2015).
[94] R. Aaij *et al.* (LHCb Collaboration), Phys. Rev. D **95**, 071101 (2017).
[95] J. McDonald, Phys. Rev. Lett. **88**, 091304 (2002).
[96] L. J. Hall, K. Jedamzik, J. March-Russell, and S. M. West, J. High Energy Phys. 03 (2010) 080.
[97] F. Elahi, C. Kolda, and J. Unwin, J. High Energy Phys. 03 (2015) 048.
[98] M. A. G. Garcia, Y. Mambrini, K. A. Olive, and M. Peloso, Phys. Rev. D **96**, 103510 (2017).
[99] I. A. Zelko, T. Treu, K. N. Abazajian, D. Gilman, A. J. Benson, S. Birrer, A. M. Nierenberg, and A. Kusenko, Phys. Rev. Lett. **129**, 191301 (2022).
[100] M. Tanabashi *et al.* (Particle Data Group Collaboration), Phys. Rev. D **98**, 030001 (2018).
[101] P. F. Depta, M. Hufnagel, and K. Schmidt-Hoberg, J. Cosmol. Astropart. Phys. 05 (2020) 009.
[102] J. Jaeckel, P. C. Malta, and J. Redondo, Phys. Rev. D **98**, 055032 (2018).
[103] S. Hoof and L. Schulz, J. Cosmol. Astropart. Phys. 03 (2023) 054.
[104] E. Müller, F. Calore, P. Carenza, C. Eckner, and M. C. D. Marsh, J. Cosmol. Astropart. Phys. 07 (2023) 056.

[105] M. Diamond, D. F. G. Fiorillo, G. Marques-Tavares, and E. Vitagliano, Phys. Rev. D **107**, 103029 (2023); **108**, 049902(E) (2023).
[106] A. Caputo, H.-T. Janka, G. Raffelt, and E. Vitagliano, Phys. Rev. Lett. **128**, 221103 (2022).
[107] M. Diamond, D. F. G. Fiorillo, G. Marques-Tavares, I. Tamborra, and E. Vitagliano, Phys. Rev. Lett. **132**, 101004 (2024).
[108] C. O'Hare, cajohare/axionlimits: Axionlimits, https://cajohare.github.io/AxionLimits/ (2020).
[109] E. W. Kolb and S. Wolfram, Phys. Lett. **91B**, 217 (1980).
[110] E. W. Kolb and S. Wolfram, Nucl. Phys. **B172**, 224 (1980); **B195**, 542(E) (1982).
[111] J. N. Fry, K. A. Olive, and M. S. Turner, Phys. Rev. D **22**, 2953 (1980).
[112] V. Shtabovenko, R. Mertig, and F. Orellana, Comput. Phys. Commun. **256**, 107478 (2020).
[113] P. Gondolo and G. Gelmini, Nucl. Phys. **B360**, 145 (1991).