

Constructing a light QCD axion without $1/N$ tuning

Abhishek Banerjee^{*,†}, Manuel A. Buen-Abad[‡], and Anson Hook[‡]

Maryland Center for Fundamental Physics, University of Maryland, College Park, Maryland 20742, USA



(Received 11 August 2025; accepted 18 September 2025; published 15 October 2025)

We consider axions lighter than what their QCD couplings might otherwise suggest. Starting with a $\mathbb{Z}_N$ axion, we introduce a small explicit $\mathbb{Z}_N$ symmetry-breaking coupling between the Standard Model Higgs boson and a reheaton. This small explicit breaking allows us to populate a large portion of the light axion $m_a - f_a$ plane, removes the $1/N$ tuning in the $\mathbb{Z}_N$ -axion, and explains why only our sector was reheated. Because of finite temperature effects, axions of this sort undergo either “rigged” misalignment, where the axion misalignment angle is effectively π regardless of its initial value, or “shuffled” misalignment, where the initial angle is effectively randomized.

DOI: [10.1103/xdyq-b67n](https://doi.org/10.1103/xdyq-b67n)

I. INTRODUCTION

One of the intriguing properties of the Standard Model (SM) is the smallness of the neutron electric dipole moment (EDM). This puzzling fact, sometimes called the strong CP problem, begs for an explanation. A very popular solution to this problem is the QCD axion [1–8]. The QCD axion dynamically relaxes the neutron EDM to zero, thereby addressing the strong CP problem; at the same time, the axion can be cold dark matter (DM), economically solving two problems at the same time [9–11].

The QCD axion is highly predictive. The axion mass is uniquely determined by its coupling to QCD, f_a , which yields the well-known QCD axion line. As a popular DM candidate, there are a large number of experiments looking for the QCD axion (see, e.g., Refs. [12–14] and references therein for a comprehensive list). For a fixed mass, m_a , any given experiment will first exclude couplings stronger than the QCD axion before, hopefully, working its way to the QCD axion line itself. Because the QCD axion has a fixed QCD coupling for a given mass, it is interesting to ask if the experiments are excluding *any* theory on their way to the QCD axion line. For couplings to fermions and photons, the answer to this question is that the experiments are probing a wide variety of axionlike particle (ALP) models.

The situation is different for the defining coupling of the QCD axion, because the coupling to gluons is also the

coupling responsible for producing the mass of the axion. Experimentally, this feature leads to a variety of interesting signatures, as finite density effects can be very important [15–26]. Theoretically, to get an axion with the same mass but larger coupling or, equivalently, the same coupling but a smaller mass, it becomes necessary to cancel the calculable QCD contribution to the mass of the axion. There is only one theory that achieves this without fine-tuning: the $\mathbb{Z}_N$ axion [27].

The $\mathbb{Z}_N$ axion invokes several decoupled copies of the Standard Model (SM). By having the axion nonlinearly realize a $\mathbb{Z}_N$ exchange symmetry, there is an exponential cancellation between the various contributions to the axion mass, making it much lighter than in the canonical QCD case.

There are, however, several features and/or issues with the $\mathbb{Z}_N$ axion. Because N is an integer, the theory cannot accommodate every single combination of small m_a and f_a , but only a discrete set of them. This is good news for falsifiability, as an axion discovery with $m_a - f_a$ values outside of this discrete set of combinations would instantly disprove the $\mathbb{Z}_N$ model, but is also a problem as this implies that there is *no* theory that naturally populates the entire parameter space.

The next issue is that observations require that only one of the identical N SM copies has a significant cosmic abundance; all other sectors must have a negligible energy density. Observational constraints on $\Delta N_{\text{eff}} \equiv \rho_{\text{dark rad.}}/\rho_{1\nu} \lesssim 0.3$ [28] imply that the temperature of any of the $\mathbb{Z}_N$ copies of the SM must be much smaller than that of our own, spontaneously breaking the $\mathbb{Z}_N$ symmetry. In other words, $\mathbb{Z}_N$ -breaking dynamics *must* be part of our cosmic history. Finally, the last issue to address is that, as a consequence of the nonlinear realization of the $\mathbb{Z}_N$ symmetry, the $\mathbb{Z}_N$ axion has N minima, only one of which solves the strong CP problem for our sector (“our” sector

* Contact author: abanerj4@umd.edu

† Contact author: buenabad@umd.edu

‡ Contact author: hook@umd.edu

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

being defined as the sector that was reheated). Thus, solving our strong CP problem requires the axion to land in the correct minimum, which constitutes a $1/N$ accident.

In this paper, we seek to address all of the previous issues with the $\mathbb{Z}_N$ axion. Our approach is to introduce an explicit, but very small, soft $\mathbb{Z}_N$ symmetry-breaking coupling between the SM Higgs and the reheaton, the particle¹ whose decay reheats our sector. This explicit breaking allows our copy of the SM to be reheated while all the other copies remain cold. Because of this soft symmetry-breaking coupling, our Higgs boson has a minutely different mass from that of the other Higgses copies, resulting in an additional contribution to the axion mass that is aligned with $\bar{\theta} = 0$ for our sector. With the appropriate sign, this additional contribution removes the N vacua degeneracy in favor of a single vacua with $\bar{\theta} = 0$, solving the $1/N$ tuning problem. Finally, because there is an additional, extremely small parameter in the reheaton-Higgs potential, the axion mass m_a and f_a are decoupled from each other. As a result, most of the light axion region in the $m_a - f_a$ plane can be populated. Finally, the global shape of the axion potential is just the canonical QCD axion potential rescaled to smaller masses. The distinction between a $\mathbb{Z}_N$ axion potential and re-scaled QCD axion potential is important when it comes to neutron star constraints, as it is not clear if the bounds of Refs. [23,24] apply to the $\mathbb{Z}_N$ axion, but do apply to the model we present in this paper.

The cosmology of our light axions is slightly different from that of the $\mathbb{Z}_N$ axion [29,30]. As we discuss below, there are two main regions of the axion parameter space, which correspond to two different misalignment regimes. In the first regime, which we dub “rigged misalignment,” there is an early period of oscillations around $\theta = \pi$, effectively “rigging” the axion initial conditions to this specific value, regardless of what they might have been before. As a result of this rigging of the axion’s initial conditions, for this parameter space region, the correct dark matter abundance is obtained for a fixed line. In the second regime, which we call “shuffled misalignment,” there is also an early era of oscillations around $\theta = \pi$; however, this is followed by a period during which most of the axion energy is kinetic, resulting in a randomized effective axion misalignment angle. This randomization is sensitive to percent-level changes in the initial condition, and thus it can be said that it “reshuffles” them “fairly.” Much like in the canonical misalignment mechanism scenario, our shuffled misalignment regime can result in axion dark matter having the correct abundance, granted some luck in its initial conditions.

In the rest of this paper, we flesh out this idea in more detail. In Sec. II, we discuss the model. In Sec. III, we discuss the cosmological evolution of our model. In Sec. III B, we discuss its implications for dark matter. Finally, we conclude in Sec. IV.

II. MODEL

In this section we present our model, which is a modification of the $\mathbb{Z}_N$ axion. We begin with a review of the original model, after which we provide a description of our changes, as well as their effects on the axion potential.

A. $\mathbb{Z}_N$ -axion

The starting point of the $\mathbb{Z}_N$ axion is N identical copies of the SM, related to each other by a $\mathbb{Z}_N$ exchange symmetry that is also nonlinearly realized on the axion [27]:

$$\mathbb{Z}_N: \mathcal{L}_{\text{SM},k} \rightarrow \mathcal{L}_{\text{SM},k+1} \quad \text{and} \quad \frac{a}{f_a} \rightarrow \frac{a}{f_a} + \frac{2\pi}{N}. \quad (2.1)$$

The Lagrangian terms relevant for the axion are

$$\mathcal{L} \supset \frac{\alpha_s}{8\pi} \sum_{k=0}^{N-1} \left(\frac{a}{f_a} + \bar{\theta} + \frac{2\pi k}{N} \right) (G\tilde{G})_k, \quad (2.2)$$

where α_s is the strong coupling constant. The k th SM copy has a strong CP angle given by $\bar{\theta}_k \equiv \bar{\theta} + 2\pi k/N$, and without loss of generality, we take the $k=0$ SM copy to be our sector.

The axion coupling to each of the $(G\tilde{G})_k$ generates an axion potential. To leading order in chiral perturbation theory (χ PT), the canonical QCD axion potential generated by this coupling is

$$V(\theta + \bar{\theta}) = -\frac{m_\pi^2 f_\pi^2}{1+z} \sqrt{1+z^2+2z \cos(\theta + \bar{\theta})}, \quad (2.3)$$

where we define $\theta \equiv a/f_a$ and $m_\pi \approx 135$ MeV, $f_\pi \approx 93$ MeV, and $z = m_u/m_d = 0.47$ are the pion mass, pion decay constant, and the up-to-down quark mass ratio, respectively [14]. This potential gives the QCD axion a mass,

$$(m_a^2)_{\text{QCD}} = \frac{z}{(1+z)^2} \frac{m_\pi^2 f_\pi^2}{f_a^2}. \quad (2.4)$$

The total axion potential generated from the couplings in Eq. (2.2) can be written as

$$V_N(\theta) = \sum_{k=0}^{N-1} V(\theta + \bar{\theta}_k) \quad (2.5)$$

$$\begin{aligned} &\xrightarrow{N \gg 1} \frac{m_\pi^2 f_\pi^2}{\sqrt{N\pi}} \sqrt{\frac{1-z}{1+z}} (-z)^N \cos N\theta + \mathcal{O}(z^{2N}) \\ &\equiv (-1)^N \frac{(m_a^2)_N f_a^2}{N^2} \cos N\theta. \end{aligned} \quad (2.6)$$

For notational convenience, we have used a shift in the axion field in order to set $\bar{\theta} = 0$, and we have included the

¹This could be the inflaton itself, or another particle.

large N limit of the potential [31]. Furthermore, we have dropped all constant terms (i.e., those not depending on the axion field itself), and in the last equality, we ignored the $\mathcal{O}(z^{2N})$ terms. Therefore, for large N , the axion mass is suppressed with respect to the canonical case, namely,

$$\epsilon_N \equiv \frac{(m_a^2)_N}{(m_a^2)_{\text{QCD}}} \simeq \sqrt{\frac{1-z^2}{\pi}} (1+z) N^{3/2} z^{N-1}. \quad (2.7)$$

This exponential suppression is due to the analytic nature of the axion potential as well as to the fact that Eq. (2.5) is a Riemann sum [27,31].

B. Our modification

As mentioned before, we introduce a soft, $\mathbb{Z}_N$ symmetry-breaking coupling between the reheaton, denoted by ϕ , and the Higgs, H . Namely, we consider the coupling

$$\mathcal{L} \supset -\phi \sum_{k=0}^{N-1} g_k |H_k|^2 - V_\phi(\phi), \quad (2.8)$$

where g_k is a dimension-1 coupling. A coupling of this form can be enforced by requiring an additional $\mathbb{Z}_2$ symmetry under which both ϕ and g_k are odd.

For $m_\phi \gtrsim 2m_h$, the reheaton decays into a pair of Higgses H_k , thereby reheating the k th SM sector. If all of the g_k were equal, then the $\mathbb{Z}_N$ symmetry would be restored, and all sectors would be equally reheated. For simplicity, let us assume maximal $\mathbb{Z}_N$ symmetry breaking and, without loss of generality, choose $g \equiv g_0 \neq 0$ and $g_{k \neq 0} = 0$. This guarantees that, upon decaying, the reheaton reheats only one sector: ours.²

1. Higgs vacuum expectation value

Aside from causing asymmetric reheating, Eq. (2.8) also induces a small shift to the Higgs vacuum expectation value (VEV).

To see exactly how this works, let us begin by writing down the potential for the SM ($k=0$) Higgs and the ϕ reheaton:

$$V(\phi, H) = -(\mu^2 - g\phi)(H^\dagger H) + \lambda_H (H^\dagger H)^2 + \frac{1}{2} m_\phi^2 \phi^2 \pm g \frac{\Lambda_{\text{NP}}^2}{16\pi^2} \phi. \quad (2.9)$$

The last term is an estimate for the size of the linear term allowed by symmetries. This estimate arises from closing

²We would like to reiterate that what makes our sector *ours* is the fact that it is the only one currently present in the Universe. Thus, for a maximally $\mathbb{Z}_N$ -breaking reheaton-Higgs sector (where the reheaton couples to only one of the Higgses), we can always call the reheated sector ours. Choosing $k=0$ as its label is entirely arbitrary.

the Higgs loop in the ϕ self-energy diagram, and it involves a UV cutoff Λ_{NP} where the effective field theory breaks down. Throughout this paper, we take g to be the smallest scale in the problem, $g \ll m_\phi \lesssim \Lambda_{\text{NP}}$. A typical value of g is $g \sim 10^{-7} m_\phi$. This choice of g is technically natural since, in the limit of $g \rightarrow 0$, the higher-spin symmetries for ϕ [32] are restored. We have neglected the ϕ^3 , ϕ^4 , and $\phi^2 |H|^2$ couplings in our analysis, as for our parameter choices these interactions are small. Their expected sizes are $\gtrsim \mathcal{O}(g^3/16\pi^2 m_h^2)$, $\mathcal{O}(g^4/16\pi^2 m_h^4)$, and $\mathcal{O}(g^2/m_h^2)$, respectively.

After electroweak symmetry breaking, both ϕ and H obtain VEVs, which can be written as

$$\begin{aligned} \langle \phi \rangle &= \phi_0 = -\frac{g}{m_\phi^2} \left(\frac{v_0^2}{2} \pm \frac{\Lambda_{\text{NP}}^2}{16\pi^2} \right), \\ \langle H^2 \rangle &= \frac{v_0^2}{2} \simeq \frac{v_{k \neq 0}^2}{2} \left[1 + \frac{g^2}{2\lambda_H m_\phi^2} \left(1 \pm \frac{\Lambda_{\text{NP}}^2}{8\pi^2 v_{k \neq 0}^2} \right) \right], \end{aligned} \quad (2.10)$$

where $v_0 = 246$ GeV is the VEV of our sector's Higgs, and $v_{k \neq 0}$ as the VEV of the other $H_{k \neq 0}$ Higgses. We have replaced $\mu^2/\lambda_H = v_{k \neq 0}^2$ and, in the last equality, performed a Taylor expansion for $g^2/2\lambda_H m_\phi^2 \ll 1$. There is a runaway behavior of the potential for $\lambda_H \lesssim g^2/(2m_\phi^2)$, which can be cured by introducing higher-order terms such as $\phi^2 |H|^2$, ϕ^4 , etc. For our purposes, we take instead $\lambda_H \gg g^2/(2m_\phi^2)$ and obtain from Eq. (2.10)

$$\frac{\Delta v}{v_0} \equiv \frac{v_0 - v_{k \neq 0}}{v_0} \simeq \frac{g^2}{4\lambda_H m_\phi^2} \left(1 \pm \frac{\Lambda_{\text{NP}}^2}{8\pi^2 v_0^2} \right). \quad (2.11)$$

Hence, we find that, due to the interaction between the SM Higgs and the reheaton, the Higgs VEV of our sector is changed with respect to every other $N-1$ sector with $k \neq 0$. As will become clear in the subsequent sections, for our model to solve the strong CP problem, we need the VEV of our sector's Higgs to be larger than the VEV of the other Higgses; therefore, we choose $\pm \rightarrow +$.

The end result of this modification to the $\mathbb{Z}_N$ axion is a standard Higgs portal, which has been extensively studied in the literature (see, e.g., Refs. [33–35] and references therein). While we do not review the Higgs portal scenario in this paper, we want to highlight an important point. Demanding that one-loop radiative corrections be no larger than the tree-level values for m_ϕ^2 , m_h^2 , and λ_H , we obtain a constraint on g as $g/m_\phi \lesssim \min[4\pi m_h/m_\phi, \sqrt{\lambda_H}]$ [ignoring the cutoff-dependent logarithmic factor, which is $\sim \mathcal{O}(10)$ for $\Lambda_{\text{NP}} = 100$ TeV]. This constraint is easily satisfied for the sufficiently small values of g in which we are interested.

2. Axion potential

When the Higgs VEVs of the N sectors are identical, the axion potential is given by Eq. (2.5). When the Higgs VEV of our $k = 0$ sector is changed, the axion potential is changed to

$$\begin{aligned}\bar{V}(\theta) &= (1 + \epsilon_\phi)V(\theta) + \sum_{k=1}^{N-1} V\left(\theta + \frac{2\pi k}{N}\right) \\ &= \epsilon_\phi V(\theta) + V_N(\theta).\end{aligned}\quad (2.12)$$

$\epsilon_\phi \sim \Delta v/v_0$ comes from a Taylor series of Eq. (2.3) with a small change in the Higgs VEV and is explicitly derived in Eq. (2.14). Around $\theta = 0$, the axion mass can be written as

$$\frac{m_a^2}{(m_a^2)_{\text{QCD}}} = (-1)^{N+1}\epsilon_N + \epsilon_\phi \ll 1. \quad (2.13)$$

Thus, unlike in the original $\mathbb{Z}_N$ model, as long as $1 \gg \epsilon_\phi \gtrsim \epsilon_N$, our model yields a light QCD axion *for any N* (even or odd). Since the ϕ -induced contribution to the potential is only minimized at $\theta = 0$, the strong CP problem of our sector is *always* solved, without the need for a linear $1/N$ tuning. Finally, as $m_a \simeq \sqrt{\epsilon_\phi}(m_a)_{\text{QCD}}$, in our model, the axion cannot be made arbitrarily light, unlike in the $\mathbb{Z}_N$ case. As will be clearer when cosmology is discussed in the next section, ϵ_ϕ is connected to the reheating temperature, and thus the minimum ϵ_ϕ our model permits is $\mathcal{O}(10^{-24})$ (for $T_{\text{rh}} \sim 10$ MeV). Thus, the axion in our model can be up-to $\mathcal{O}(10^{-12})$ lighter than the canonical QCD axion (see Fig. 1).

The last step in our axion potential calculation is to derive ϵ_ϕ . The change in the Higgs VEV shifts both the light quark masses and the confinement scale (through the running of α_s). Using the one-loop β function, we find that the change in the QCD scale due to the shift in the Higgs VEV as $\Lambda_{\text{QCD},0}/\Lambda_{\text{QCD},k \neq 0} = (v_0/v_{k \neq 0})^{2/9}$ for three light quark flavors [36]. From Eq. (2.3) and using $m_\pi^2 f_\pi^2 \propto m_q \Lambda_{\text{QCD}}^3 \propto v^{5/3}$ with $q = u, d$, we find

$$\begin{aligned}\epsilon_\phi &\equiv \frac{\Delta(m_\pi^2 f_\pi^2)}{m_\pi^2 f_\pi^2} = \frac{5}{3} \frac{\Delta v}{v_0} \\ &\simeq \frac{5g^2}{12\lambda_H m_\phi^2} \left(1 + \frac{\Lambda_{\text{NP}}^2}{8\pi^2 v_0^2}\right).\end{aligned}\quad (2.14)$$

We also use Eq. (2.11) in deriving the above formula.

III. AXION COSMOLOGY

In this section, we discuss the cosmological implications of our model. We begin by analyzing the reheating phase and show how the presence of a light QCD axion is intertwined with the reheating temperature. Subsequently,




FIG. 1. Plot of the lightness of the axion mass, i.e., $\sqrt{\epsilon_\phi}$ as a function of the reheating temperature T_{rh} , for different choices of Λ_{NP} and m_ϕ . The red and turquoise lines show the parameter space for $\Lambda_{\text{NP}} = 100$ TeV and 3 TeV, respectively. The solid lines indicate the minimum value of ϵ_ϕ that can be obtained for a given Λ_{NP} with $m_\phi/\Lambda_{\text{NP}} = 1$, whereas the dashed red and dashed turquoise lines are for $m_\phi = 10$ and 0.5 TeV, respectively. The values of ϵ_ϕ in the turquoise shaded excluded region require $m_\phi \gtrsim \Lambda_{\text{NP}}$ for $\Lambda_{\text{NP}} = 3$ TeV, where the validity of the EFT breaks down. The sudden drop in the values of ϵ_ϕ around $T_{\text{rh}} \sim 0.2$ GeV is due to the change in $g_*(T)$ close to the QCD phase transition.

we investigate the viability of axion dark matter in our model.

A. Thermal history

To obtain the reheating temperature of our sector, we calculate the reheaton decay width to the SM Higgs $\phi \rightarrow H^\dagger H$ as

$$\Gamma_\phi \simeq \frac{g^2}{8\pi m_\phi}, \quad (3.1)$$

in the high-energy limit. For simplicity, we will set $(m_\phi)_{\text{min}} \sim 500$ GeV as the smallest ϕ mass for which the reheaton can still efficiently decay into a Higgs pair and thereby reheat the Universe.

Our sector is reheated to a temperature T_{rh} defined as $\Gamma_\phi = H(T_{\text{rh}}) \equiv H_{\text{rh}} = \sqrt{\pi^2 g_*(T_{\text{rh}})/90} T_{\text{rh}}^2/M_{\text{Pl}}$ (with $M_{\text{Pl}} = 2.4 \times 10^{18}$ GeV being the reduced Planck mass), while the other $N - 1$ sectors remain cold. Thus, any ΔN_{eff} constraints are readily avoided by the simple fact that the ϕ reheaton decays only to the $k = 0$ sector. The reheating temperature is therefore

$$\frac{T_{\text{rh}}}{1 \text{ TeV}} \simeq \left(\frac{g/m_\phi}{2 \times 10^{-7}}\right) \sqrt{\frac{m_\phi}{1 \text{ TeV}}} \left(\frac{106.75}{g_*(T_{\text{rh}})}\right)^{1/4}. \quad (3.2)$$

Other decay channels of ϕ to the SM are subdominant as they are suppressed by the $h - \phi$ mixing angle,

$\sin^2 \theta_{h\phi} \sim (gv_0/m_\phi^2)^2 \ll 1$. Since the coupling that shifts the Higgs VEV also leads to reheating, by combining Eqs. (3.2) and (2.14), we obtain

$$\epsilon_\phi \simeq \frac{10\pi}{3\lambda_H} \left(1 + \frac{\Lambda_{\text{NP}}^2}{8\pi^2 v_0^2}\right) \left(\frac{H_{\text{rh}}}{m_\phi}\right). \quad (3.3)$$

Given that $\epsilon_\phi \propto T_{\text{rh}}^2$, a smaller value of ϵ_ϕ is tied to a lower reheating temperature. To emphasize this, let us estimate the smallest ϵ_ϕ can be.

From Eq. (3.3), we have $\epsilon_\phi \propto \Lambda_{\text{NP}}^2 T_{\text{rh}}^2 / m_\phi$. Therefore, in order to obtain the smallest ϵ_ϕ , we need Λ_{NP} and T_{rh} to be as small as possible, while on the other hand keeping m_ϕ as large as possible. We conservatively take $(\Lambda_{\text{NP}})_{\text{min}} = 4\pi v_0 \approx 3$ TeV. The largest allowed ϕ mass, for which the effective field theory (EFT) description does not break down, is of order $\sim \mathcal{O}(\Lambda_{\text{NP}})$. Thus, taking $\Lambda_{\text{NP}} = m_\phi = 3$ TeV and a very low reheating temperature of $T_{\text{rh}} = 10$ MeV, we find that the smallest value of ϵ_ϕ allowed by our model is $\epsilon_\phi \approx 3 \times 10^{-24}$. For comparison, another reasonable benchmark point in parameter space is $\Lambda_{\text{NP}} = 100$ TeV, $m_\phi = 10$ TeV, and $T_{\text{rh}} = 1$ TeV, which yields the larger value of $\epsilon_\phi \approx 2 \times 10^{-11}$. In Fig. 1, we plot this minimum ϵ_ϕ for a given cutoff of the theory, as a function of T_{rh} . The turquoise and the red lines represent $\Lambda_{\text{NP}} = 3$ and 100 TeV, respectively, with the solid and dashed textures corresponding to different choices of m_ϕ . The sudden jumps in ϵ_ϕ are caused by the abrupt change in the degrees of freedom g_* of the thermal bath around the QCD phase transition.

For the other $N - 1$ sectors not to thermalize with the SM bath via a tree-level a exchange [processes such as $(G\tilde{G})_{k=0} \rightarrow (G\tilde{G})_{k \neq 0}$], its associated rate must satisfy $\Gamma_{0 \rightarrow k} \sim T^5 (\alpha_s / 8\pi f_a)^4 \lesssim H_{\text{rh}}$, resulting in the following bound:

$$f_a \gtrsim 18 \text{ TeV} \left(\frac{\alpha_s}{0.09}\right) \left(\frac{T_{\text{rh}}}{1 \text{ TeV}}\right)^{3/4}. \quad (3.4)$$

More conservatively, we impose an even stronger condition by demanding that the axion not be thermalized [e.g., via $(G\tilde{G})_{k=0} \rightarrow aG_{k=0}$]. The rate associated with this process must obey $\Gamma_{0 \rightarrow a} \sim T^3 (g_s \alpha_s / 8\pi f_a)^2 \lesssim H_{\text{rh}}$, which results in [37]

$$f_a \gtrsim 4 \times 10^8 \text{ GeV} \left(\frac{\alpha_s}{0.09}\right)^{3/2} \left(\frac{T_{\text{rh}}}{1 \text{ TeV}}\right)^{1/2}. \quad (3.5)$$

For $T_{\text{rh}} \lesssim T_c$, we consider axion thermalization through the pions and nucleons of the $k = 0$ th sector (e.g., via $\pi\pi \rightarrow a\pi$ and/or $aN \rightarrow N\pi$ processes). Neglecting phase space factors, which only loosen the constraint, Ref. [38] obtains

$$f_a \gtrsim 2 \times 10^7 \text{ GeV} \left(\frac{T}{0.1 \text{ GeV}}\right)^{3/2} \quad (3.6)$$

for the axion not to be thermalized. Throughout this paper, we take f_a such that Eqs. (3.4) and (3.5) or Eq. (3.6) is satisfied.

B. Axion dark matter

In this subsection, we explore the axion cosmology of our model. We will pay special attention to the regions of parameter space where the axion can also play the role of dark matter. While all of our plots are obtained numerically, the majority of this section will instead be dedicated to obtaining an analytical understanding of the dynamics.

Since only the $k = 0$ sector is reheated, the temperature (T) dependence of the axion potential originates from the 0th sector. The temperature-dependent axion potential can be written as

$$\begin{aligned} \bar{V}(T, \theta) &\simeq (1 + \epsilon_\phi) F(T) V(\theta) + \sum_{k=1}^{N-1} V\left(\theta + \frac{2\pi k}{N}\right) \\ &= V(\theta) [(1 + \epsilon_\phi) F(T) - 1] + V_N(\theta), \end{aligned} \quad (3.7)$$

where $F(T)$ captures the temperature dependence of the topological susceptibility. Roughly speaking, $F(T)$ has the behavior $F(T) = (T_c/T)^8$ for $T \gg T_c$, and $F(T) = 1$ for $T \ll T_c$, with $T_c \simeq 150$ MeV being the QCD phase transition temperature [39–41].³ The axion potential can be written explicitly as

$$\begin{aligned} \frac{\bar{V}(T, \theta)}{(m_a^2)_{\text{QCD}} f_a^2} &= (-1)^N \frac{\epsilon_N}{N^2} \cos(N\theta) \\ &\quad + \frac{(1+z)}{z} [1 - (1 + \epsilon_\phi) F(T)] \\ &\quad \times \sqrt{1 + z^2 + 2z \cos \theta} + \text{const.} \end{aligned} \quad (3.8)$$

At high temperatures, $F(T) \ll 1$, and the above potential is minimized at $\theta = \pi$. The axion mass around this high-temperature minimum can be written as $m_a^\pi \equiv m_a(T \gg T_c) \simeq (m_a)_{\text{QCD}} \sqrt{(1+z)/(1-z)}$, which for $z \approx 0.5$ is larger than the canonical QCD axion by a factor of $\sim \sqrt{3}$. On the other hand, for $T \ll T_c$, $F(T) \simeq 1$, and the potential is minimized at $\theta = 0$, with the axion mass suppressed by ϵ_ϕ , as discussed previously. Therefore,

³The exact functional form of $F(T)$ close to the phase transition temperature is an active area of research [39–41]. In our numerical calculations, we interpolate between the high-temperature and low-temperature behaviors by writing $F(T) = 1/[1 + (T/T_c)^{8\beta}]^{1/\beta}$, for $\beta > 0$. A sharper (smooth) phase transition can be achieved for $\beta \gg 1$ ($\beta \ll 1$). For our qualitative estimates, the exact value of β bears no significant change. In all our plots below, we take $\beta = 12$.

unlike in the case of the canonical QCD axion, in our model, the axion was parametrically several orders of magnitude *heavier* in the early Universe than it is today (by a factor of about $\sqrt{3/\epsilon_\phi}$ for $\epsilon_\phi \ll \epsilon_N$); additionally, in the early Universe, the minimum of the high-temperature potential was located at $\theta = \pi$ instead of $\theta = 0$. This represents a significant departure from the standard misalignment mechanism, where the potential is flat at very high temperatures and a significant tuning is required to obtain an initial angle $\sim \pi$ (cf. Refs. [29,31]).

It is now evident that, for large reheating temperature $T_{\text{rh}} \gtrsim T_c$, the axion in our model can undergo a two-stage misalignment mechanism: first evolving toward the high-temperature minimum at $\theta = \pi$, before eventually turning around and evolving toward the low-temperature global minimum at $\theta = 0$. Since $m_a^\pi \approx \sqrt{3}(m_a)_{\text{QCD}}$, the two oscillatory stages occur as long as $(m_a)_{\text{QCD}} \gtrsim H(T_c) \equiv H_c$, irrespective of the value of ϵ_ϕ . Numerically, we find that the axion undergoes one full oscillation around this minimum if $f_a \lesssim (f_a)_{\text{crit}} \approx 5 \times 10^{16}$ GeV, corresponding to $(m_a)_{\text{QCD}} \approx 8H(T_c)$. The time $t_{\text{osc},\pi}$, at which this first era of oscillations begins, is obtained by solving $3H(t_{\text{osc},\pi}) = \min[m_a^\pi, 3H_{\text{rh}}]$, and the corresponding amplitude of the axion field can be written as

$$\pi - \theta(t) \simeq |\pi - \theta_i| \left(\frac{a_{\text{osc},\pi}}{a} \right)^{3/2} \cos[m_a^\pi(t - t_{\text{osc},\pi})], \quad (3.9)$$

where $\theta_i \equiv \theta(t_i)$ is the initial misalignment angle of the axion field at some initial time t_i , which we take to be the time of reheating $t_i = t_{\text{rh}}$. Thus, we have $\theta_i = \theta_{\text{rh}}$.⁴ As T nears the QCD phase transition temperature, $F(T) \rightarrow 1$, and the temperature-dependent part of the axion potential vanishes. The potential develops a new minimum at $\theta = 0$, furthermore becoming shallower as its amplitude gets multiplied by ϵ_ϕ . At this point, the axion mass changes from $m_a^\pi \approx \sqrt{3}(m_a)_{\text{QCD}}$ to $(m_a)_{\text{QCD}}\sqrt{\epsilon_\phi} \ll (m_a)_{\text{QCD}}$.⁵ At the time t_c of the QCD phase transition, the axion will be very close to π and will have some velocity leftover from the first era of oscillations, parametrically given by $|\dot{\theta}_c| \sim m_a^\pi |\pi - \theta_c|$.

The next stage of axion evolution depends on how its zero-temperature mass compares to the Hubble scale at T_c .

⁴Henceforth, we denote $\theta_x \equiv \theta(t_x)$ the value of the axion field $\theta(t)$ at some time of interest t_x . The same is true of the axion field velocity, namely, $\dot{\theta}_x \equiv \dot{\theta}(t_x)$.

⁵The axion potential, and thus the axion mass, vanishes when $T_c/T = [(1 + \epsilon_\phi)^\beta - 1]^{-1/8\beta} \simeq (\beta\epsilon_\phi)^{-1/8\beta}$. Since $(\beta\epsilon_\phi)^{-1/8\beta} \gtrsim 1$ for $\epsilon_\phi \lesssim 1$, the amplitude of the axion potential starts to decrease as $T \rightarrow T_c$, and it vanishes at some temperature just after the phase transition, which is dictated by β . The next phase of the axion evolution depends on the zero-temperature axion mass, i.e., $m_a = (m_a)_{\text{QCD}}\sqrt{\epsilon_\phi}$, and on $H(T_c)$. We find for $\beta = 12$ the temperature at which axion mass vanishes is ~ 100 MeV for $\epsilon_\phi = 10^{-20}$, and it gets closer to T_c as we increase ϵ_ϕ .

In the first case, which we call “shuffled misalignment,” $H_c \gtrsim m_a = (m_a)_{\text{QCD}}\sqrt{\epsilon_\phi}$. During shuffled misalignment, the axion field value is effectively randomized before it starts oscillating again. In the second case, which we call “rigged misalignment,” $m_a = (m_a)_{\text{QCD}}\sqrt{\epsilon_\phi} \gtrsim H_c$. During rigged misalignment, the axion field value settles close $\theta \approx \pi$ irrespective of the initial field value, and its subsequent evolution is determined by its leftover velocity.

1. Shuffled misalignment

Let us first focus on the $H_c \gtrsim m_a = (m_a)_{\text{QCD}}\sqrt{\epsilon_\phi}$ case, for which the axion potential is extremely flat around the phase transition. Because of this, the axion is kinetic energy dominated so that $\ddot{\theta} + 3H\dot{\theta} \simeq 0 \Rightarrow \dot{\theta} \propto a^{-3}$. The kinetic energy-dominated phase continues until $H \sim (m_a)_{\text{QCD}}\sqrt{\epsilon_\phi}$, at which point the axion starts to oscillate around the CP -preserving minimum. Thus, in between the QCD phase transition and the onset of the second oscillatory stage, the field excursion $\Delta\theta \equiv \theta_{\text{osc},0} - \theta_c$ can be written as $\Delta\theta \sim \pm|\pi - \theta_{\text{rh}}|((m_a)_{\text{QCD}}/H_c)^{1/4}$, where $\theta_{\text{osc},0} = \theta(t_{\text{osc},0})$ is the field amplitude at the onset of the second oscillatory era around $\theta = 0$, and $3H(T_{\text{osc},0}) = (m_a)_{\text{QCD}}\sqrt{\epsilon_\phi}$. The choice in $\pm$ in the above equation depends on the sign of $\dot{\theta}_c$, which is effectively random due to the fast axion oscillations around $\theta = \pi$. Because $(m_a)_{\text{QCD}}, H_{\text{rh}} \gg H_c$, $\Delta\theta$ can easily be of order $\mathcal{O}(\pi)$. Thus, the rolling phase generically makes $\theta_{\text{osc},0}$ very different from θ_{rh} , effectively randomizing (shuffling) the misalignment angle in preparation for the second oscillatory phase. Numerically, we find that initial conditions differing by at least $\mathcal{O}(1\%)$ are enough to change $\theta_{\text{osc},0}$ by order 1.

Once the Hubble scale becomes comparable to $(m_a)_{\text{QCD}}\sqrt{\epsilon_\phi}$, the field starts to oscillate once again, but this time around $\theta = 0$. The energy density of the axion can be written as $\rho_a(t) = m_a^2 f_a^2 \theta^2(t)/2$, which can account for present day dark matter energy density when

$$f_a \simeq 2 \times 10^{14} \text{ GeV} \left(\frac{\epsilon_\phi}{2 \times 10^{-10}} \right)^{-1/6} \left(\frac{\theta_{\text{osc},0}}{\pi/\sqrt{3}} \right)^{-4/3} \left(\frac{r}{1} \right)^{2/3}, \quad (3.10)$$

where $r = \Omega_a/\Omega_{\text{DM}}$ and $\pi/\sqrt{3}$ serves as a benchmark for the rms value of an initial amplitude for θ (in the harmonic limit) [42].

In the above equation, we have taken $\theta_{\text{osc},0}$ to be a free parameter, as it depends sensitively on all of the other parameters in the theory. As stated above, even percent-level changes in θ_{rh} can change $\theta_{\text{osc},0}$ by an order-1 amount. Numerically, this extreme sensitivity comes about at the transition between the two phases, namely, when the axion oscillates around $\theta = \pi$ and when the axion becomes kinetic energy dominated. Because even two very similar initial conditions may result in $\mathcal{O}(1)$ differences in their




FIG. 2. *Shuffled misalignment.* The evolution of the axion angle θ (top) and its normalized energy density $\rho_a t_c^2 / f_a^2$ (bottom) as a function of time t , for $f_a \simeq 2 \times 10^{15}$ GeV and $\epsilon_\phi = 10^{-7}$. Top: the evolution of $\theta(t)$ for three very similar initial conditions: $\theta_{\text{rh}} = \pi/4$ (blue) and $\theta_{\text{rh}} = \pi/4(1 \pm 10\%)$ (red and green, respectively). In yellow, we plot the evolution of a canonical QCD axion with same f_a and $\theta_{\text{rh}} = \pi/4$. The dashed (solid) black line shows $t_{\text{osc},\pi}$ ($t_{\text{osc},0}$), when the axion begins to oscillate around $\theta = \pi$ ($\theta = 0$). The gray dashed vertical line marks t_c . Note that, due to the periodicity of the axion potential, the vertical axis in the top panel must be taken to be modulo 2π . Therefore, the final oscillatory stage, shown to be around $\theta = 2\pi$, can equivalently be interpreted to take place around $\theta = 0$. Bottom: the evolution of the energy density of the axion. The solid lines correspond to the total energy density, while the thin lines correspond to the kinetic energy term only. The colors have the same meaning as in the top panel, and so do the vertical lines.

respective axion field values after the QCD phase transition, we deem this scenario “shuffled misalignment.”

In Fig. 2, we depict the evolution of the axion field (θ , top panel) and its normalized energy density ($\rho_a t_c^2 / f_a^2$, bottom panel) in the shuffled misalignment regime as a function of time t , for $f_a = 2 \times 10^{15}$ GeV and $\epsilon_\phi = 10^{-7}$. In the top panel, the three lines plot $\theta(t)$ for three very similar initial conditions: $\theta_{\text{rh}} = \pi/4$ (blue) and $\theta_{\text{rh}} = \pi/4(1 \pm 10\%)$ (red and green, respectively). For comparison, we show as well the evolution of a canonical QCD axion with the same f_a , and $\theta_{\text{rh}} = \pi/4$, in yellow. There are clearly four distinct periods of evolution for the axion. In the first, the axion value remains frozen at the initial value θ_{rh} . Around the time $t = t_{\text{osc},\pi}$ (black dashed

vertical line, and as discussed above), the axion begins to oscillate around $\theta = \pi$, due to important temperature effects in its potential. The QCD phase transition occurs at t_c (gray dashed vertical line), at which point the temperature dependence of the axion potential quickly vanishes. In this third epoch, the axion energy density becomes dominated by its kinetic contribution, and the axion field rolls away from π . At this point, even very similar initial conditions will result drastically different evolutions for θ , depending on the field velocity $\dot{\theta}_c$. Finally, at $t_{\text{osc},0}$ (black solid vertical line), the axion begins its second stage of oscillations, this time around $\theta = 2\pi$ (or equivalently, $\theta = 0$), with a misalignment angle effectively given by $\theta_{\text{osc},0}$. As discussed above, the preceding rolling era effectively shuffles the value of the axion angle, which can result in vastly different $\theta_{\text{osc},0}$ values even for very similar initial conditions. In the bottom panel, we show the evolution of the energy density of the axion, normalized by t_c^2 / f_a^2 . The thick lines depict the total energy density, while the thin lines show only the kinetic term. The blue, red, and green colors correspond to the same initial conditions as in the top panel. In this panel, we can easily see how, despite starting their evolution with very similar values, the axion energy density for each initial condition drastically changes soon after the QCD phase transition, due to the shuffling of the misalignment angle. During the rolling phase, between the dashed gray and solid black vertical lines, the kinetic energy accounts for most of the axion energy density.

2. Rigged misalignment

Let us now consider the opposite case and study the evolution of axions heavier than H_c , i.e., $m_a = \sqrt{\epsilon_\phi}(m_a)_{\text{QCD}} \gtrsim H_c \approx 10^{-11}$ eV. In this case, the axion begins to oscillation around the $\theta = 0$ minimum immediately after the QCD phase transition, with initial amplitude $\theta_c \sim \pi$ and a small velocity $|\dot{\theta}_c| \sim m_a^\pi |\pi - \theta_c|$. Thus, unlike in the previous case, the field undergoes both oscillatory phases one right after the other, without an intervening rolling phase. In the absence of the rolling phase, the initial amplitude for the second phase of the oscillation is rigged to be $\sim \pi$ irrespective of the initial field value. This behavior is why we called this “rigged misalignment.”

As a result of rigging, the axion abundance in the late-time Universe can simply be obtained as

$$\rho_a(t \gg t_c) \simeq \frac{\pi^2}{2} \epsilon_\phi (m_a)_{\text{QCD}}^2 f_a^2 (t_c/t)^{3/2}, \quad (3.11)$$

which matches the DM energy density for

$$\epsilon_\phi \simeq 7.8 \times 10^{-8} (r/1). \quad (3.12)$$

Notice that for a given ϵ_ϕ the estimate in Eq. (3.11) is independent of f_a [since $(m_a)_{\text{QCD}} f_a = \text{const.}$], as long as




FIG. 3. *Rigged misalignment*. the evolution of the axion angle θ (top) and its normalized energy density $\rho_a t_c^2 / f_a^2$ (bottom) as a function of time t , for $f_a \approx 2 \times 10^{15}$ GeV and $\epsilon_\phi = 0.1$. Top: the evolution of $\theta(t)$ for three very different initial conditions: $\theta_{\text{th}} = \pi/4$ (solid blue), $\theta_{\text{th}} = \pi/2$ (solid red), and $\theta_{\text{th}} = 3\pi/4$ (dashed green). The vertical and horizontal lines have the same meaning as those in Fig. 2. Note that, due to the periodicity of the axion potential, the vertical axis in the top panel must be taken to be modulo 2π . Therefore, the blue and green curves' oscillations around $\theta = 2\pi$ are the same as those of the red curve around $\theta = 0$. Bottom: the evolution of the energy density of the axion shifted by a constant such that when $t < t_c$, $\theta = \pi$ has zero energy and when $t > t_c$, $\theta = 0$ has zero energy. The colors have the same meaning as in the top panel, and so do the vertical lines. The black horizontal line is $\rho(\theta = \pi) - \rho(\theta = 0)$. The fact that, regardless of initial conditions, the axion total energy density at $t = t_c$ jumps to the black line demonstrates that it is restarting the misalignment mechanism with an angle rigged to $\theta = \pi$.

$(m_a)_{\text{QCD}} \gtrsim (m_a)_{\text{QCD}} \sqrt{\epsilon_\phi} \gtrsim H_c$, i.e., for $f_a \lesssim 2 \times 10^{14}$ GeV $(\epsilon_\phi / 7.8 \times 10^{-8})^{1/2}$.

In Fig. 3, we depict the evolution of the axion field (θ , top panel) and its normalized energy density ($\rho_a t_c^2 / f_a^2$, bottom panel) in the rigged misalignment regime, with $f_a = 2 \times 10^{15}$ GeV and $\epsilon_\phi = 0.1$. The black dashed and solid vertical lines have the same meaning as those in Fig. 2. In the top panel, the three lines correspond to $\theta(t)$ with very different initial conditions: $\theta_{\text{th}} = \pi/4$ (solid blue), $\theta_{\text{th}} = \pi/2$ (solid red), and $\theta_{\text{th}} = 3\pi/4$ (dashed




FIG. 4. The $m_a - f_a$ parameter space of our model, for the $T_{\text{rh}} \gtrsim T_c$ case. Here, we have taken $T_{\text{rh}} = 1$ GeV. The dashed green and pink horizontal lines correspond to $f_a = (f_a)_{\text{crit}}$ and $f_a = M_{\text{Pl}}$, respectively. The blue line corresponds to the canonical QCD axion. The brown region is ruled out by various astrophysical and cosmological bounds. The dashed green vertical line corresponds to $m_a = H_c$, which represents the boundary between the shuffled and rigged regimes of the two-stage misalignment mechanism of our model. Since $\epsilon_\phi < 1$, both regimes lie to the left of the blue line. In the purple region, $\epsilon_\phi \lesssim (\epsilon_\phi)_{\text{min}}$ as given by Eq. (3.3). In this plot, we get $(\epsilon_\phi)_{\text{min}} \approx 9 \times 10^{-20}$, for $m_\phi = \Lambda_{\text{NP}} = 3$ TeV, and $T_{\text{rh}} = 1$ GeV. The black line corresponds to those values of f_a and m_a for which the axion can account for the entirety of the dark matter. In the shuffled regime, one can still get the right dark matter abundance above or below this line by tuning $\theta_{\text{osc},0}$ in the ways indicated in the plots. In the rigged regime, all points above (below) the black line, shown by the light red shading, correspond to an underabundance (overabundance) of axion dark matter. We have omitted the experimental bounds on this parameter space to avoid overcrowding the plot; they can be seen in Fig. 5.

green). As in Fig. 2, the yellow line plots the canonical QCD axion, with same f_a and $\theta_{\text{th}} = \pi - 10^{-3}$. The two oscillatory eras described above, the first around $\theta = \pi$ and beginning at $t_{\text{osc},\pi}$ and the second around $\theta = 0$ (or equivalently 2π) and beginning at $t_{\text{osc},0} = t_c$, can be easily seen in this figure. Unlike in the shuffled misalignment case, however, the rolling phase is not present, and the two oscillatory phases take place in quick succession after the QCD phase transition. Because of this, the effective misalignment angle for the final oscillatory stage is $\theta_{\text{osc},0} = \theta_c \approx \pi$. An initial misalignment angle so close to π typically delays the onset of oscillations due to the flatness of the potential, as can be seen in the yellow line (see also Refs. [43,44]). However, since in our model there is a nonzero field velocity at t_c , there is no such delay, and the axion field begins its second phase of oscillations immediately after the phase transition.

Since the effective misalignment angle of the last phase has been rigged to be $\approx \pi$ regardless of the initial conditions θ_{th} , the subsequent evolution of the axion field and,




FIG. 5. The parameter space of our model when $T_{\text{rh}} \lesssim T_c$. In this case, there is no first stage of oscillations around $\theta = \pi$, and thus there is neither a shuffled nor a rigged regime. The dashed pink and blue lines and brown region have the same meaning as those in Fig. 4. Since $\epsilon_\phi < 1$, our model lies to the left of the blue line. The dashed green vertical lines correspond to $m_a = 3H_{\text{rh}}$, for $T_{\text{rh}} = 10$ MeV and $T_{\text{rh}} = 100$ MeV. These lines represent the boundary between when the axion oscillates some time afterward reheating (left of the line) or the more model-dependent case where it starts oscillating before or at reheating (right of the line). In the purple region, $\epsilon_\phi \lesssim (\epsilon_\phi)_{\text{min}}$ as given by Eq. (3.3). For this plot, $(\epsilon_\phi)_{\text{min}} \approx 3 \times 10^{-24}$, for $m_\phi = \Lambda_{\text{NP}} = 3$ TeV, and $T_{\text{rh}} = 10$ MeV. The black line corresponds to those values of f_a and m_a for which the axion can account for the entirety of the dark matter. One can still get the right dark matter abundance above or below this line depending on the value of θ_{rh} . The gray shaded regions represent combined exclusions from various model dependent astrophysical, cosmological, and laboratory experimental results [15, 23–25, 46–68].

consequently, of its energy density, is the same. This can be seen by the overlap between the solid blue and dashed green lines. The same is true of the red line which, despite appearing to be very different from the green and blue ones, is actually almost identical to them due to the mod 2π nature of the θ angle.

In the bottom panel, we once again show the evolution of the energy density of the axion, normalized by t_c^2/f_a^2 . The blue, red, and green colors correspond to the same initial conditions as in the top panel, while the black horizontal line shows $\rho(\theta = \pi) - \rho(\theta = 0)$. In this panel, we can easily see how, despite starting their evolution with very different values, the axion energy densities for each initial condition all converge to the same evolution soon after the QCD phase transition, due to the rigged nature of the misalignment angle. This can be explicitly seen when $t = t_c$, where the axion begins oscillating around $\theta = 0$ with a rigged effective initial misalignment angle of π .

3. Folded misalignment

So far, we have seen that for $T_{\text{rh}} \gtrsim T_c$ and $f_a \lesssim (f_a)_{\text{crit}}$ the axion undergoes two oscillatory eras, whether in the

shuffled or rigged regimes. However, it is possible for the axion to “fold,” or not partake in the dynamics of this first stage for sufficiently large f_a ($f_a \gtrsim (f_a)_{\text{crit}} \approx 5 \times 10^{16}$ GeV) and/or a small reheating temperature ($T_{\text{rh}} \lesssim T_c \sim 150$ MeV). The axion field will always eventually oscillate around $\theta = 0$ with an initial amplitude of $\theta_{\text{osc},0} \simeq \theta_{\text{rh}}$ as long as m_a is larger than the Hubble expansion rate today. The oscillation time is obtained by solving $3H(T_{\text{osc},0}) = \min[3H_{\text{rh}}, m_a]$ like the case of standard ALP dark matter. Of course, when $m_a > 3H_{\text{rh}}$, θ_{rh} will depend heavily on the dynamics preceding reheating. This case is similar to that of the standard misalignment for ALPs, and as such, we defer a detail discussion of it to Appendix B.

Let us conclude this section with a quick discussion about how requiring our axion to account for the totality of dark matter influences the parameter space. Perhaps the most conservative demand we can make of an axion dark matter scenario is that it begins to oscillate no later than when the Universe had a temperature of $T_{\text{osc},0} \sim \text{keV}$, which corresponds to the time when the smallest modes probed by current cosmic microwave background observations entered the horizon [45]. A quick calculation shows that such a constraint translates into a bound of $\epsilon_\phi \gtrsim 4 \times 10^{-21}$ for $f_a = M_{\text{pl}}$, not very far from the independent bound on ϵ_ϕ outlined in Sec. III A.

IV. CONCLUSION

In Figs. 4 and 5, we show our model’s $m_a - f_a$ parameter space, for $T_{\text{rh}} \gtrsim T_c$ and $T_{\text{rh}} \lesssim T_c$, respectively. In Fig. 4, we took $T_{\text{rh}} = 1$ GeV; its precise value has no impact on the late-time abundance of the axion as long as it is above T_c . On the other hand, we allow T_{rh} to vary between 10 and 100 MeV in Fig. 5. In both figures, the dashed green and pink horizontal lines correspond to $f_a = (f_a)_{\text{crit}}$ and $f_a = M_{\text{pl}}$, respectively. The canonical QCD axion is shown as a blue line [3–8]. Since $\epsilon_\phi \lesssim 1$, our model lives to the left of this blue line. The gray shaded regions represent combined exclusions from various model dependent astrophysical, cosmological, and laboratory experimental results [15, 23–25, 46–68]. To avoid clutter, these constraints are only shown in Fig. 5 but not shown in Fig. 4; however, the constraints apply equally well in both scenarios. In purple, we shade the region for which ϵ_ϕ is so small that the desired reheat temperature cannot be obtained [$\epsilon_\phi \lesssim (\epsilon_\phi)_{\text{min}}$ as given by Eq. (3.3)]. Taking $m_\phi = \Lambda_{\text{NP}} = 3$ TeV, we get $(\epsilon_\phi)_{\text{min}} \approx 9 \times 10^{-20}$ for $T_{\text{rh}} = 1$ GeV in Fig. 4 and $(\epsilon_\phi)_{\text{min}} \approx 3 \times 10^{-24}$ for $T_{\text{rh}} = 10$ MeV in Fig. 5. In Fig. 4, the dashed green vertical line corresponds to $m_a = H_c \approx 10^{-11}$ eV, which roughly separates shuffled and rigged misalignment regimes.

In the rigged regime, the black line indicates when the axion can account for the entirety of the dark matter. Equation (3.11) tells us that, since $\theta_{\text{osc},0} \approx \pi$, the axion dark

matter will be underabundant or overabundant above or below this black line. In the shuffled regime, the black line indicates when the axion can account for the entirety of the dark matter for a typical misalignment angle of $\theta_{\text{osc},0} = \pi/\sqrt{3}$. Equation (3.10) indicates that the correct DM abundance can still be obtained above or below the black line, by choosing initial conditions such that $\theta_{\text{osc},0}$ lies closer to $\theta = 0$ or to $\theta = \pi$, respectively.

In Fig. 5, as $T_{\text{rh}} \lesssim T_c$, the axion skips the first stage of oscillations around $\theta = \pi$, and its evolution is like that of a light ALP. The dashed green lines correspond to the transition between the axion oscillating some time after reheating (to the left of the vertical lines) to oscillating sometime before reheating takes place (to the right of the vertical lines) for $T_{\text{rh}} = 10$ and 100 MeV. To avoid discussing the details of reheating, we present our plots in terms of θ_{rh} with the understanding that its value will depend on the unspecified dynamics occurring before reheating. As in the previous figure, the black line corresponds to points in the parameter space for which the axion dark matter has the observed abundance for a typical misalignment angle of $\theta_{\text{rh}} \approx \pi/\sqrt{3}$; tuning this angle closer to 0 or π allows the axion to be the dark matter above or below this line.

In this work, we demonstrated that by introducing a soft $\mathbb{Z}_N$ -breaking coupling, one can simultaneously populate the light QCD-axion parameter space up to masses $\mathcal{O}(10^{-12})$ times lighter than their canonical QCD values, while resolving the usual $1/N$ tuning and reheating ambiguities of the original $\mathbb{Z}_N$ QCD axion model. We extended the $\mathbb{Z}_N$ QCD axion framework by introducing a $\mathbb{Z}_N$ -symmetry breaking coupling, g , between the Higgs and the reheaton field (ϕ), through which the reheaton decays exclusively into our Standard Model sector. This coupling also induces a slight shift in the Higgs VEV in our sector compared to the others. As a result, the sector that gets reheated has a slightly larger axion potential, leading to an imperfect $\mathbb{Z}_N$ cancellation that selects the $\theta = 0$ minimum.

Because of the shifted Higgs VEV, the axion mass becomes $m_a^2 \simeq (\epsilon_\phi + \epsilon_N)(m_a)_{\text{QCD}}^2$, where ϵ_N is the axion mass in the $\mathbb{Z}_N$ limit. The presence of $\epsilon_\phi \sim g^2/m_\phi^2$ decouples the axion mass from its decay constant f_a , allowing the model to accommodate a lighter axion for any $\epsilon_N \lesssim \epsilon_\phi \ll 1$. Thus, unlike the original $\mathbb{Z}_N$ construction (which predicts a discrete set of axion masses for odd values of N), our setup continuously populates the axion parameter space for arbitrary N .

Cosmologically, this setup gives rise to three distinct misalignment regimes depending on the reheating temperature and the axion decay constant. For high reheating temperatures ($T_{\text{rh}} \gtrsim \text{GeV}$) and small $f_a \lesssim 5 \times 10^{16}$ GeV, early oscillations around $\theta = \pi$ pin the axion near π until the QCD phase transition. For heavier axions with $m_a \gtrsim H_c \approx 10^{-11}$ eV, the field begins oscillating around

the $\theta = 0$ minimum immediately after the QCD phase transition. This dynamically fixes the late-time misalignment angle—and therefore the relic abundance—along a single trajectory in the (m_a, f_a) plane, regardless of the initial condition at the time of reheating. We refer to this regime as “rigged” misalignment. In contrast, for $m_a \lesssim H_c \approx 10^{-11}$ eV, the axion undergoes a rolling phase after initially oscillating around π . This rolling randomizes the effective misalignment angle, allowing for the usual ALP-like tuning to reproduce the observed dark matter energy density. We refer to this as the “shuffled” regime. For $f_a \gg 5 \times 10^{16}$ GeV and/or low reheating temperatures, the axion skips the initial oscillations around $\theta = \pi$, resulting in the standard ALP dark matter scenario. We refer to this as “folded” misalignment.

In this work, we do not construct a UV-complete model that simultaneously generates the $\mathbb{Z}_N$ symmetric coupling at the scale f_a and naturally explains a small symmetry-breaking Higgs-reheaton coupling at a lower scale. While such symmetry-breaking features can be readily embedded within a KSVZ-type UV completion of the $\mathbb{Z}_N$ axion, we defer a detailed exploration of this possibility to future work. Also, currently in our model, we focus on providing analytical estimates for various dark matter scenarios without performing full numerical simulations. In particular, in the rigged misalignment regimes, where the axion begins with an initial misalignment angle of $\mathcal{O}(\pi)$, a detailed understanding of dark matter perturbation growth and structure formation may require incorporating nonlinear dynamics and topological effects, which lie beyond the scope of this study. Finally, a thorough analysis of the full temperature-dependent axion potential during a kination phase could unveil novel cosmological signatures. We leave this investigation to future work as well.

ACKNOWLEDGMENTS

The authors would like to thank Sebastian A. R. Ellis for providing us with tabulated data of the existing constraints. The authors would also like to thank Michael Geller and Konstantin Springmann for useful discussions and reading the draft. The authors are supported by the National Science Foundation under Grant No. PHY2210361 and the Maryland Center for Fundamental Physics.

DATA AVAILABILITY

The data that support the findings of this article are openly available (<https://github.com/papai1994/Light-QCD-axion>).

APPENDIX A: QUALITY PROBLEM, $\mathbb{Z}_N$, AND LIGHT QCD AXION

We devote this appendix to the status of the quality problem in our model. In the framework of the EFT, one

generically expects that the global Peccei-Quinn (PQ) symmetry, which gives rise to the axion, is broken by quantum gravity effects [69–71], although with the expectation that such effects decouple in the limit $M_{\text{Pl}} \rightarrow \infty$.

Let us begin by discussing the quality problem of the canonical QCD axion. In this case, one can, for instance, write operators of the form

$$V_\Delta = \left(\frac{c_p \Phi^{p+4}}{M_{\text{Pl}}^p} + \text{H.c.} \right) + (\text{similar-order terms}) \\ \rightarrow \left(\frac{f_a}{M_{\text{Pl}}} \right)^p f_a^4 \cos(p\theta + \beta), \quad (\text{A1})$$

which explicitly breaks the PQ symmetry and generates a potential for the QCD axion. Here, $\beta = \text{Arg}(c_p)$ is a generic phase, which is set to $\mathcal{O}(1)$. In general, these Planck suppressed (“quality”) operators shift the minima of the axion potential, which leads to a constraint on p ,

$$p \gtrsim \frac{125 + 4 \log(f_a/10^{10} \text{ GeV})}{19 - \log(f_a/10^{10} \text{ GeV})}, \quad (\text{A2})$$

in order for the leading operator in Eq. (A1) not to spoil the axion solution to the strong CP problem and generate a neutron EDM.

In our mode, where the QCD axion is lighter than usual, one might naively expect that the quality of the axion is worsened, since the discrete $\mathbb{Z}_N$ symmetry is broken softly by g . However, since the underlying $\mathbb{Z}_N$ symmetry is restored as $g \rightarrow 0$, it can be gauged, which means that the leading-order operators spoiling the axion quality should depend on g . More importantly, since $(g, \phi) \rightarrow (-g, -\phi)$ is still a symmetry of the system, these leading-order operators should be $\propto g^2$. Therefore, in our model, the leading PQ symmetry-breaking operator takes the form of

$$V_\Delta = g^2 \left(\frac{c_p \Phi^{p+2}}{M_{\text{Pl}}^p} + \text{H.c.} \right) + (\text{similar-order terms}) \\ \rightarrow \left(\frac{f_a}{M_{\text{Pl}}} \right)^p g^2 f_a^2 \cos(p\theta + \beta). \quad (\text{A3})$$

For $\epsilon_\phi \gtrsim \epsilon_N$, using Eqs. (2.12) and (2.14), we find that these operators do not run afoul neutron EDM experiments for

$$p \gtrsim \frac{93 - 2 \log[(\Lambda_{\text{NP}}/m_\phi)] + 2 \log(f_a/10^{10} \text{ GeV})}{19 - \log(f_a/10^{10} \text{ GeV})}. \quad (\text{A4})$$

In Fig. 6, we show the quality of our light axion by green lines and depict the quality of the canonical QCD axion in red. As can be seen from Eq. (A3), the value of p for a given f_a depends on $\epsilon_\phi/g^2 \propto \Lambda_{\text{NP}}^2/m_\phi^2$ for $\Lambda_{\text{NP}} \gtrsim v_{k \neq 0}$. The solid line shows the value of p for a given f_a for $m_\phi/\Lambda_{\text{NP}} = 1$,




FIG. 6. Comparison of the dimension of the quality operators for the canonical QCD axion (red) and our model of the lighter QCD axion (green). We see that for a given f_a , despite being lighter than the canonical QCD axion, the axion in our model requires less protection from the Planck suppressed operators than the canonical QCD axion case. For the solid lines, we take $\Lambda_{\text{NP}} = m_\phi = 3 \text{ TeV}$, whereas for the dashed line, we take $\Lambda_{\text{NP}} = 100 \text{ TeV}$ and $m_\phi = 10 \text{ TeV}$.

whereas the dashed line depicts for $m_\phi/\Lambda_{\text{NP}} = 0.1$. Compared to the case of a canonical QCD axion, the quality of our light axion is better, as long as $\Lambda_{\text{NP}} f_a / (v_{k \neq 0} m_\phi) \gtrsim 1$. Since $m_\phi \lesssim \Lambda_{\text{NP}}$, we obtain a better quality axion for any $f_a \gtrsim v_0$. Finally, we note that the effects of $\mathbb{Z}_N$ -symmetric, Planck-suppressed operators are small, since they are suppressed by powers of $(f_a/M_{\text{Pl}})^N$ (for a specific UV completion), which is negligible in the large- N limit in which we are interested (see, e.g., Ref. [72] for a discussion).

APPENDIX B: FOLDED MISALIGNMENT

In this section, we discuss the evolution of the axion in the folded misalignment region of parameter space. As mentioned in the main text, folded misalignment occurs for large f_a ($f_a \gtrsim (f_a)_{\text{crit}} \approx 5 \times 10^{16} \text{ GeV}$) and/or a small reheating temperature ($T_{\text{rh}} \lesssim T_c \sim 150 \text{ MeV}$). In these cases, the axion DM folds on any complicated new dynamics and is produced via the regular misalignment mechanism, similar to the case of ALPs.

a. Large $f_a \gtrsim 5 \times 10^{16} \text{ GeV}$. If the axion decay constant is large, i.e., $f_a \gg (f_a)_{\text{crit}}$, then regardless of whether or not the reheat temperature was larger than the QCD scale, the field does not undergo the first era of oscillations, since for such f_a values $H(t_c) \gg (m_a)_{\text{QCD}}$. The field remains frozen for a long time and only starts to oscillate around $\theta = 0$ at $T_{\text{osc},0}$ when $3H(T_{\text{osc},0}) = m_a = (m_a)_{\text{QCD}} \sqrt{\epsilon_\phi}$. In this case, the high-temperature potential plays no role, and $\theta_{\text{osc},0} \approx \theta_{\text{rh}}$. The axion accounts for the present-day dark matter energy density for

$$f_a \simeq 10^{17} \text{ GeV} \left(\frac{\epsilon_\phi}{8 \times 10^{-17}} \right)^{-\frac{1}{6}} \left(\frac{\theta_{\text{rh}}}{0.1} \right)^{-\frac{4}{3}} \left(\frac{r}{1} \right)^{\frac{2}{3}}. \quad (\text{B1})$$

Thus, a mild tuning of the initial condition is needed for the axion to account for 100% of the dark matter.

We finish our discussion of this scenario by noting that, for $f_a \sim (f_a)_{\text{crit}}$, the first oscillatory phase is replaced by a rolling phase, in which the field slowly evolves toward $\theta = \pi$. Here, the field reaches $\theta = \pi$ only if $(m_a)_{\text{QCD}}/H(T_c) \sim \mathcal{O}(5)$. In general, the axion will either roll straight past π or undergo at most half an oscillation before rolling past it. More importantly, the field is effectively displaced from its initial value, and the misalignment angle for the second phase of oscillation gets shuffled, i.e., ends up being very different from its original value. We already consider this case while plotting the parameter space of our model in Fig. 4.

b. $T_{\text{rh}} \lesssim T_c$. For a reheating temperature smaller than T_c , the temperature corrections to the axion potential are essentially negligible. Thus, the axion potential is always minimized at $\theta = 0$, and its mass is always given by its $T = 0$ value. Unlike the case of $T_{\text{rh}} \gtrsim T_c$, the absence of the temperature effects ensures that the axion does not undergo its earlier oscillatory and/or rolling phases prior to oscillating around $\theta = 0$. Once again, this case is similar to that of

the standard ALP dark matter; the field starts to oscillate at $t_{\text{osc},0}$ which is obtained by solving $H(T_{\text{osc},0}) = \min[H_{\text{rh}}, m_a/3]$ with an initial amplitude of $\theta_{\text{osc},0} \simeq \theta_{\text{rh}}$. When $m_a \gtrsim 3H_{\text{rh}}$, θ_{rh} becomes extremely sensitive to the unspecified cosmological history before reheating. When $H(T_{\text{osc},0}) = H_{\text{rh}}$, i.e., for large m_a , the axion abundance becomes independent of f_a , and its parametric dependence is given as $\rho_a \propto \epsilon_\phi/T_{\text{rh}}^3$. We find that the axion relic constitutes 100% of the present-day dark matter energy density for

$$\epsilon_\phi = 2 \times 10^{-8} \left(\frac{H_{\text{rh}}}{H(0.1 \text{ GeV})} \right)^{3/2} \left(\frac{1}{\theta_{\text{rh}}} \frac{r}{1} \right)^2, \quad (\text{B2})$$

for $f_a \lesssim 4.8 \times 10^{13} \text{ GeV} \sqrt{\epsilon_\phi/(2 \times 10^{-8})/(H_{\text{rh}}/0.1 \text{ GeV})}$. Since $T_{\text{rh}} \gtrsim 10 \text{ MeV} \gtrsim T_{\text{BBN}} \sim \text{MeV}$, this occurs for

$$f_a \lesssim 6 \times 10^{15} \text{ GeV} \sqrt{\frac{\epsilon_\phi}{2 \times 10^{-8}}}. \quad (\text{B3})$$

On the other hand, if $H_{\text{rh}} \gtrsim m_a/3$, the axion starts to oscillate when $3H(T_{\text{osc},0}) = m_a$, and its abundance follows the scaling given in Eq. (B1).

-
- [1] R. D. Peccei and H. R. Quinn, *Phys. Rev. D* **16**, 1791 (1977).
 - [2] R. D. Peccei and H. R. Quinn, *Phys. Rev. Lett.* **38**, 1440 (1977).
 - [3] S. Weinberg, *Phys. Rev. Lett.* **40**, 223 (1978).
 - [4] F. Wilczek, *Phys. Rev. Lett.* **40**, 279 (1978).
 - [5] A. R. Zhitnitsky, *Sov. J. Nucl. Phys.* **31**, 260 (1980).
 - [6] M. Dine, W. Fischler, and M. Srednicki, *Phys. Lett.* **104B**, 199 (1981).
 - [7] J. E. Kim, *Phys. Rev. Lett.* **43**, 103 (1979).
 - [8] M. A. Shifman, A. I. Vainshtein, and V. I. Zakharov, *Nucl. Phys.* **B166**, 493 (1980).
 - [9] J. Preskill, M. B. Wise, and F. Wilczek, *Phys. Lett.* **120B**, 127 (1983).
 - [10] M. Dine and W. Fischler, *Phys. Lett.* **120B**, 137 (1983).
 - [11] L. F. Abbott and P. Sikivie, *Phys. Lett.* **120B**, 133 (1983).
 - [12] A. Ringwald, [arXiv:2404.09036](https://arxiv.org/abs/2404.09036).
 - [13] C. O'Hare, [cajohare/axionlimits: Axionlimits](https://cajohare.github.io/AxionLimits/), <https://cajohare.github.io/AxionLimits/> (2020).
 - [14] S. Navas *et al.* (Particle Data Group), *Phys. Rev. D* **110**, 030001 (2024).
 - [15] A. Hook and J. Huang, *J. High Energy Phys.* **06** (2018) 036.
 - [16] R. Balkin, J. Serra, K. Springmann, and A. Weiler, *J. High Energy Phys.* **07** (2020) 221.
 - [17] A. Prabhu, *Phys. Rev. D* **104**, 055038 (2021).
 - [18] R. Budnik, H. Kim, O. Matsedonskyi, G. Perez, and Y. Soreq, *Phys. Rev. D* **104**, 015012 (2021).
 - [19] R. Balkin, J. Serra, K. Springmann, S. Stelzl, and A. Weiler, *J. High Energy Phys.* **06** (2022) 023.
 - [20] R. Balkin, J. Serra, K. Springmann, S. Stelzl, and A. Weiler, *SciPost Phys.* **14**, 071 (2023).
 - [21] D. Noordhuis, A. Prabhu, C. Weniger, and S. J. Witte, *Phys. Rev. X* **14**, 041015 (2024).
 - [22] R. Balkin, J. Serra, K. Springmann, S. Stelzl, and A. Weiler, *J. High Energy Phys.* **02** (2025) 141.
 - [23] A. Gómez-Bañón, K. Bartnick, K. Springmann, and J. A. Pons, *Phys. Rev. Lett.* **133**, 251002 (2024).
 - [24] M. Kumamoto, J. Huang, C. Drischler, M. Baryakhtar, and S. Reddy, *Phys. Rev. D* **112**, 043008 (2025).
 - [25] R. Balkin, J. Serra, K. Springmann, S. Stelzl, and A. Weiler, *Phys. Rev. D* **109**, 095032 (2024).
 - [26] A. Banerjee, I. M. Bloch, Q. Bonnefoy, S. A. R. Ellis, G. Perez, I. Savoray, K. Springmann, and Y. V. Stadnik, [arXiv:2502.04455](https://arxiv.org/abs/2502.04455).
 - [27] A. Hook, *Phys. Rev. Lett.* **120**, 261802 (2018).
 - [28] N. Aghanim *et al.* (Planck Collaboration), *Astron. Astrophys.* **641**, A5 (2020).
 - [29] L. Di Luzio, B. Gavela, P. Quilez, and A. Ringwald, *J. Cosmol. Astropart. Phys.* **10** (2021) 001.
 - [30] R. T. Co, T. Gherghetta, Z. Liu, and K.-F. Lyu, *J. High Energy Phys.* **09** (2024) 145.

- [31] L. Di Luzio, B. Gavela, P. Quilez, and A. Ringwald, *J. High Energy Phys.* **05** (2021) 184.
- [32] A. Hook and R. Rattazzi, *Phys. Rev. D* **108**, 115019 (2023).
- [33] D. O’Connell, M. J. Ramsey-Musolf, and M. B. Wise, *Phys. Rev. D* **75**, 037701 (2007).
- [34] F. Piazza and M. Pospelov, *Phys. Rev. D* **82**, 043533 (2010).
- [35] E. Fuchs, O. Matsedonskyi, I. Savoray, and M. Schlaffer, *J. High Energy Phys.* **04** (2021) 019.
- [36] Z. Berezhiani, L. Gianfagna, and M. Giannotti, *Phys. Lett. B* **500**, 286 (2001).
- [37] E. Masso, F. Rota, and G. Zsembinszki, *Phys. Rev. D* **66**, 023004 (2002).
- [38] S. Hannestad, A. Mirizzi, and G. Raffelt, *J. Cosmol. Astropart. Phys.* **07** (2005) 002.
- [39] S. Borsanyi *et al.*, *Nature (London)* **539**, 69 (2016).
- [40] P. Petreczky, *J. Phys. G* **39**, 093002 (2012).
- [41] G. Grilli di Cortona, E. Hardy, J. Pardo Vega, and G. Villadoro, *J. High Energy Phys.* **01** (2016) 034.
- [42] C. A. J. O’Hare, *Proc. Sci.*, COSMICWISPers (2024) 040.
- [43] A. Arvanitaki, S. Dimopoulos, M. Galanis, L. Lehner, J. O. Thompson, and K. Van Tilburg, *Phys. Rev. D* **101**, 083014 (2020).
- [44] A. Chatrchyan, C. Eröncel, M. Koschnitzke, and G. Servant, *J. Cosmol. Astropart. Phys.* **10** (2023) 068.
- [45] L. Hui, J. P. Ostriker, S. Tremaine, and E. Witten, *Phys. Rev. D* **95**, 043541 (2017).
- [46] J. Zhang, Z. Lyu, J. Huang, M. C. Johnson, L. Sagunski, M. Sakellariadou, and H. Yang, *Phys. Rev. Lett.* **127**, 161101 (2021).
- [47] I. Schulthess *et al.*, *Phys. Rev. Lett.* **129**, 191801 (2022).
- [48] K. Blum, R. T. D’Agnolo, M. Lisanti, and B. R. Safdi, *Phys. Lett. B* **737**, 30 (2014).
- [49] J. Gué, P. Wolf, and A. Hees, *arXiv:2504.00720*.
- [50] X. Zhang, A. Banerjee, M. Leyser, G. Perez, S. Schiller, D. Budker, and D. Antypas, *Phys. Rev. Lett.* **130**, 251002 (2023).
- [51] S. Karanth *et al.* (JEDI Collaboration), *Phys. Rev. X* **13**, 031004 (2023).
- [52] T. S. Roussy *et al.*, *Phys. Rev. Lett.* **126**, 171301 (2021).
- [53] C. Abel *et al.*, *Phys. Rev. X* **7**, 041034 (2017).
- [54] X. Zhang, N. Houston, and T. Li, *Phys. Rev. D* **108**, L071101 (2023).
- [55] J. Alda, C. Brogini, G. Di Carlo, L. Di Luzio, D. Piatti, S. Rigolin, and C. Toni, *Phys. Rev. D* **111**, 035022 (2025).
- [56] A. Banerjee, D. Budker, M. Filzinger, N. Huntemann, G. Paz, G. Perez, S. Porsev, and M. Safronova, *arXiv:2301.10784*.
- [57] G. Lucente, L. Mastrototaro, P. Carena, L. Di Luzio, M. Giannotti, and A. Mirizzi, *Phys. Rev. D* **105**, 123020 (2022).
- [58] K. Springmann, M. Stadlbauer, S. Stelzl, and A. Weiler, *arXiv:2410.19902*.
- [59] L. Caloni, M. Gerbino, M. Lattanzi, and L. Visinelli, *J. Cosmol. Astropart. Phys.* **09** (2022) 021.
- [60] M. Baryakhtar, M. Galanis, R. Lasenby, and O. Simon, *Phys. Rev. D* **103**, 095019 (2021).
- [61] M. J. Stott, *arXiv:2009.07206*.
- [62] C. Ünal, F. Pacucci, and A. Loeb, *J. Cosmol. Astropart. Phys.* **05** (2021) 007.
- [63] S. Hoof, D. J. E. Marsh, J. Sisk-Reynés, J. H. Matthews, and C. Reynolds, *arXiv:2406.10337*.
- [64] S. J. Witte and A. Mummery, *Phys. Rev. D* **111**, 083044 (2025).
- [65] A. Caputo and G. Raffelt, *Proc. Sci.*, COSMICWISPers (2024) 041 [*arXiv:2401.13728*].
- [66] N. Iwamoto, *Phys. Rev. Lett.* **53**, 1198 (1984).
- [67] M. Buschmann, C. Dessert, J. W. Foster, A. J. Long, and B. R. Safdi, *Phys. Rev. Lett.* **128**, 091102 (2022).
- [68] K. Springmann, M. Stadlbauer, S. Stelzl, and A. Weiler, *J. High Energy Phys.* **02** (2025) 138.
- [69] M. Kamionkowski and J. March-Russell, *Phys. Lett. B* **282**, 137 (1992).
- [70] S. M. Barr and D. Seckel, *Phys. Rev. D* **46**, 539 (1992).
- [71] S. Ghigna, M. Lusignoli, and M. Roncadelli, *Phys. Lett. B* **283**, 278 (1992).
- [72] A. Banerjee, J. Eby, and G. Perez, *Phys. Rev. D* **107**, 115011 (2023).