Dynamical dark matter.   II. An explicit model

Keith R. Dienes[1,2,3,*] and Brooks Thomas[4,†]

[1]*Physics Division, National Science Foundation, Arlington, Virginia 22230 USA*
[2]*Department of Physics, University of Maryland, College Park, Maryland 20742 USA*
[3]*Department of Physics, University of Arizona, Tucson, Arizona 85721 USA*
[4]*Department of Physics, University of Hawaii, Honolulu, Hawaii 96822 USA*

In a recent paper [K. R. Dienes and B. Thomas, Phys. Rev. D **85,** 083523 (2012).], we introduced "dynamical dark matter," a new framework for dark-matter physics, and outlined its underlying theoretical principles and phenomenological possibilities. Unlike most traditional approaches to the dark-matter problem which hypothesize the existence of one or more stable dark-matter particles, our dynamical dark-matter framework is characterized by the fact that the requirement of stability is replaced by a delicate balancing between cosmological abundances and lifetimes across a vast ensemble of individual dark-matter components. This setup therefore collectively produces a time-varying cosmological dark-matter abundance, and the different dark-matter components can interact and decay throughout the current epoch. While the goal of our previous paper was to introduce the broad theoretical aspects of this framework, the purpose of the current paper is to provide an explicit model of dynamical dark matter and demonstrate that this model satisfies all collider, astrophysical, and cosmological constraints. The results of this paper therefore constitute an "existence proof" of the phenomenological viability of our overall dynamical dark-matter framework, and demonstrate that dynamical dark matter is indeed a viable alternative to the traditional paradigm of dark-matter physics. Dynamical dark matter must therefore be considered alongside other approaches to the dark-matter problem, particularly in scenarios involving large extra dimensions or string theory in which there exist large numbers of particles which are neutral under standard-model symmetries.

I. INTRODUCTION

The nature of what constitutes the nonbaryonic dark matter in our universe remains one of the most fundamental mysteries in particle physics [1]. The most precise measurements of the relic abundance of this dark matter to date are those derived from WMAP data [2], which yield a value

$$\Omega_{\rm CDM} h^2 = 0.1131 \pm 0.0034, \qquad (1.1)$$

where $h \approx 0.72$ is the Hubble constant. Beyond this, we know very little about the properties of this dominant constituent of the matter density in our universe, save that its interactions with the fields of the standard model (SM) are extremely weak. One of the reasons why the nature of the dark matter remains so elusive is its apparent stability. Observational constraints on the lifetime τ_χ of any decaying dark-matter candidate χ are quite stringent. Indeed, for any particle with a relic abundance $\Omega_\chi \sim \Omega_{\rm CDM}$, current limits [3] from cosmic microwave background (CMB) measurements, etc., require that

$$\tau_\chi \gtrsim 10^{26} \text{ s.} \qquad (1.2)$$

For this reason, most models of the dark sector posit the existence of a single dark-matter particle (or, in the case of certain multicomponent dark-matter scenarios [4,5], a small number of such particles) which is either absolutely stable (with that stability usually conferred by some additional symmetry, such as R-parity in supersymmetric models, KK-parity [6] in universal extra dimensions [7–9], or T-parity [10] in little-Higgs theories [11]), or else sufficiently long-lived as to satisfy the bound in Eq. (1.2). Indeed, the phenomenological consequences of dark-matter decays in models with unstable dark-matter candidates [12] can be quite significant.

Recently, an alternative framework for addressing the dark-matter question has been proposed [13]. In this so-called "dynamical dark matter" paradigm, the dark sector comprises not one or merely a few particle species, but rather a vast ensemble of different fields ϕ_i, each of which contributes only a fraction Ω_i of the total dark-matter relic abundance $\Omega_{\rm CDM}$. None of these fields is presumed to be absolutely stable, and thus a nonzero decay width Γ_i is associated with each field. However, in this framework, the individual relic abundances of the ϕ_i fields are presumed to be generated in such a way that the most stable members of that ensemble are the most abundant. By contrast, the abundances of the more unstable members are suppressed according to the size of their decay widths. It is this balancing between Γ_i and Ω_i which makes it possible for

*dienes@physics.arizona.edu
†thomasbd@phys.hawaii.edu

083524-1

the phenomenological constraints relating to the effects of dark-matter decays to be satisfied.

In Ref. [13], we focused on the model-independent aspects of our dynamical dark-matter framework, discussing its broad theoretical properties, without any detailed phenomenological analysis or comparison with data. By contrast, in this work, as a "proof of concept," we provide an explicit model of dynamical dark matter. In this model, the fields which collectively constitute the dynamical dark-matter ensemble are the KK excitations of a light axionlike field propagating in the bulk of a spacetime with one or more large, flat extra dimensions. In this model, the fields of the SM, as well as the gauge fields associated with some additional, non-Abelian gauge group G which confines at a scale Λ_G, are taken to be localized on a four-dimensional subspace of that bulk. The axion field is assumed to couple to the gauge fields of G (and also potentially to one or more of the SM fields) via nonrenormalizable operators suppressed by some effective, four-dimensional cutoff scale $\hat{f}_X$. We shall demonstrate that within this setup, the resulting ensemble of axion KK modes naturally satisfies all applicable observational constraints on dark-matter decays—even if the stability of this "dark tower" is entirely unprotected.

Another advantage of this particular dynamical dark-matter model is that it is not only phenomenologically viable, but also theoretically well motivated. In any theory in which the SM fields reside on a brane, the KK excitations of any bulk field are, from the point of view of the four-dimensional theory, massive particles neutral under the SM gauge group. Thus, were it not for the lack of a stabilizing symmetry, any of these particles would be a natural candidate for dark matter. However, our results demonstrate that a lack of stability is not an insurmountable impediment to such fields serving as dark matter *collectively*, rather than individually. This is an exciting prospect, for it provides a novel way of addressing the dark-matter question in theories with extra dimensions. Furthermore, our model also demonstrates that realizing a viable dynamical dark-matter ensemble does not require an overly complicated dark sector, a large number of independent mass scales, or an excessive degree of fine-tuning. Indeed, the model presented here involves only three independent physical scales: the effective four-dimensional cutoff scale $\hat{f}_X$, the confinement scale Λ_G, and the compactification scale M_c. Together, these three scales determine the mass spectrum and decay properties of the entire ensemble.

The outline of this paper is as follows. In Sec. II, we briefly review the formalism for discussing axions and axionlike fields, beginning with the standard, four-dimensional case and then moving on to the generalized, five-dimensional bulk-axion case. In Sec. III, we calculate the decay widths of the KK modes of such a bulk axionlike field and investigate how these decay widths scale with

the mass of the mode. In the process, we show that the decays of the lighter modes to SM fields experience a natural suppression, but that such decays nevertheless dominate over decays to other, lighter bulk fields in the theory. In Sec. IV, we examine the various mechanisms through which a population of axion modes may be generated in the early universe, and demonstrate that the abundances generated for those modes by misalignment production are indeed balanced against their decay widths in precisely the manner required for dynamical dark matter. In Sec. V, we examine the collective properties of the ensemble of axion KK modes. We show that such an ensemble can collectively reproduce the observed dark-matter relic density given in Eq. (1.1), and that it possesses the appropriate equation of state to be regarded as dark matter. In Sec. VI, we summarize the experimental, astrophysical, and cosmological constraints on scenarios involving bulk axions in large, flat extra dimensions. We demonstrate that these constraints can be satisfied in a model which also simultaneously yields the correct total relic abundance—in other words, that our model truly constitutes a viable model of dynamical dark matter. Finally, in the Conclusions, we summarize the results of the previous sections and discuss several further directions for future investigation.

As we have indicated, this paper is the second part of a two-part series that began with Ref. [13]. Consequently, we shall assume that the reader is familiar with the ideas, notation, and conventions established in Ref. [13] in what follows.

II. BULK AXIONS AS DYNAMICAL DARK MATTER

As discussed in the Introduction, the model for dynamical dark matter that we shall consider in this paper is a model in which the KK excitations of a bulk axion constitute the dark-matter ensemble. In this section, therefore, we briefly review the formalism relevant for describing the dynamics of axions in four or more dimensions. We begin with a brief summary of the relevant properties of the four-dimensional QCD axion (more detailed reviews of which may be found, e.g., in Refs. [14–17]), and then discuss how this formalism can be generalized to a broader class of axions and axionlike fields. Finally, we summarize the formalism for embedding such fields in the bulk in theories with extra dimensions.

A. Axions in four dimensions

The QCD axion emerges as a consequence of the Peccei-Quinn (PQ) mechanism [18], a mechanism which provides an elegant, dynamical solution to the strong CP problem. The strong CP problem arises due to the nontrivial vacuum structure of QCD. Specifically, the QCD Lagrangian can in principle contain an additional term

$$\mathcal{L}_{\text{QCD}} \ni \bar{\Theta} \frac{g_s^2 \xi}{32\pi^2} G^{\mu\nu a} \tilde{G}_{\mu\nu}^a, \tag{2.1}$$

where g_s is the $SU(3)$ coupling, ξ is a numerical factor of $\mathcal{O}(1)$, $G_{\mu\nu}^a$ is the field-strength tensor for the gluon field, and $\tilde{G}_{\mu\nu}^a = \frac{1}{2}\epsilon_{\mu\nu\rho\sigma}G^{\rho\sigma a}$ is its dual. The parameter $\bar{\Theta}$ is given by $\bar{\Theta} \equiv \Theta + \text{Arg det}M$, where Θ is the strong-interaction theta-angle and M is the Cabbibo-Kobayashi-Maskawa (CKM) matrix. In principle, $\bar{\Theta}$ can take any value. However, experimental bounds on the electric dipole moment d_n of the neutron serve to constrain $\bar{\Theta}$. The most stringent limit is currently $|d_n| \leq 2.9 \times 10^{-26}$ e cm [19], which translates into a bound

$$\bar{\Theta} < 0.7 \times 10^{-11}. \tag{2.2}$$

While there is, in principle, no problem with $\bar{\Theta}$ taking so small a value, there is no particular reason why it should be so small. This fine-tuning issue is what is commonly referred to as the strong CP problem.

In the Peccei-Quinn solution to the strong CP problem, the effective $\bar{\Theta}$-parameter associated with the gluon field relaxes to zero dynamically as a consequence of the spontaneous breaking of an anomalous, global $U(1)$ symmetry, usually dubbed $U(1)_{\text{PQ}}$, at some high scale f_{PQ}. The spontaneous breaking of this $U(1)_{\text{PQ}}$ symmetry implies the presence of a pseudo-Nambu-Goldstone boson: a real pseudoscalar commonly known as the QCD axion [20], which necessarily interacts with the gluon field via the Lagrangian

$$\mathcal{L} \ni \frac{1}{2}\partial_\mu a \partial^\mu a + \frac{g_s^2 \xi}{32\pi^2 f_{\text{PQ}}} a G^{\mu\nu a}\tilde{G}_{\mu\nu}^a, \tag{2.3}$$

where a denotes the axion field. Furthermore, this pseudoscalar may also have interactions with the other fields of

the SM. The presence of the anomalous $U(1)_{\text{PQ}}$ symmetry in the high-scale theory determines the effective Lagrangian for these interactions (at leading order) to be

$$\mathcal{L}_{\text{int}} = \frac{g_s^2 \xi}{32\pi^2 f_{\text{PQ}}} a G_{\mu\nu}^a \tilde{G}^{a\mu\nu} + \sum_i \frac{c_i}{f_{\text{PQ}}}(\partial_\mu a)\bar{\psi}_i\gamma^\mu\gamma^5\psi_i$$
$$+ \frac{e^2 c_\gamma}{32\pi^2 f_{\text{PQ}}} a F_{\mu\nu}\tilde{F}^{\mu\nu}, \tag{2.4}$$

where ψ_i are the SM fermions and c_γ and c_i are dimensionless coefficients. These coefficients depend on the charge assignments of the SM fields (and potentially of additional fields in the theory as well) under $U(1)_{\text{PQ}}$, and are therefore substantially more model-dependent than ξ.

At high temperatures, the axion field is effectively massless, as befits a Nambu-Goldstone boson; however, it acquires a small, temperature-dependent mass $m_a(T)$ at lower scales due to QCD instanton effects. A number of computations of this mass have been performed, and while the results depend to some extent on the assumptions and calculational techniques involved, $m_a(T)$ is often assumed to have the rough form [21,22]

$$m_a(T) \approx \frac{g_s \xi}{4\sqrt{2}\pi} \frac{\Lambda_{\text{QCD}}^2}{f_{\text{PQ}}} \times \begin{cases} b\left(\frac{\Lambda_{\text{QCD}}}{T}\right)^4 & \text{for } T \gtrsim \Lambda_{\text{QCD}} \\ 1 & \text{for } T \lesssim \Lambda_{\text{QCD}}, \end{cases} \tag{2.5}$$

where $\Lambda_{\text{QCD}} \approx 250$ MeV is the QCD confinement scale, b is a numerical coefficient of $\mathcal{O}(10^{-2})$, and ξ is an $\mathcal{O}(1)$ numerical factor.

The fact that the axion necessarily couples to the gluon field implies that it will also have effective couplings to hadrons. The most important such couplings, phenomenologically speaking, are those of the axion to pions and nucleons. These couplings take the form [23]

$$\mathcal{L}_{\text{had}} = \frac{C_{a\pi}}{f_\pi f_{\text{PQ}}}(\partial_\mu a)[(\partial^\mu \pi^+)\pi^-\pi^0 + (\partial^\mu\pi^-)\pi^+\pi^0 - 2(\partial^\mu\pi^0)\pi^+\pi^-] + \frac{C_{an}}{f_{\text{PQ}}}(\partial_\mu a)\bar{n}\gamma^\mu\gamma^5 n + \frac{C_{ap}}{f_{\text{PQ}}}(\partial_\mu a)\bar{p}\gamma^\mu\gamma^5 p$$
$$+ \frac{iC_{a\pi N}}{f_\pi f_{\text{PQ}}}(\partial_\mu a)[\pi^+(\bar{p}\gamma^\mu n) - \pi^-(\bar{n}\gamma^\mu p)]. \tag{2.6}$$

The precise values for the effective nucleon-nucleon-axion couplings C_{ap} and C_{an}, the nucleon-pion-axion coupling $C_{a\pi N}$, and the axion-pion-pion coupling $C_{a\pi}$ depend on the $U(1)_{\text{PQ}}$ charges of the quark fields. For the case of a so-called hadronic axion [24], which does not couple directly to the SM quarks, the coefficients for the axion-nucleon-nucleon interaction are

$$C_{ap} = 0.24\frac{z}{(1+z)} + 0.15\frac{z-2}{(1+z)} + 0.02,$$
$$C_{an} = 0.24\frac{z}{(1+z)} + 0.15\frac{1-2z}{(1+z)} + 0.02, \tag{2.7}$$

where $z = m_u/m_d \approx 0.56$ denotes the ratio of the up-quark to down-quark masses. Similarly, the coefficients for the interactions involving pions are given by

$$C_{a\pi N} = \frac{1-z}{2\sqrt{2}(1+z)}, \qquad C_{a\pi} = \frac{1-z}{3(1+z)}, \tag{2.8}$$

where $m_\pi \approx 135.0$ MeV is the neutral pion mass, and $f_\pi \approx 93$ MeV is the pion decay constant. These hadronic couplings play an important role in constraining the parameter space of axion models.

The QCD axion is the prototypical example of a light pseudoscalar field whose mass arises solely due to nonperturbative effects associated with instanton dynamics,

and whose interactions with the SM fields are highly suppressed. It is by no means the only example, however. Indeed, a wide variety of additional particles possessing these same properties have appeared in the literature in a number of beyond-the-Standard-Model (BSM) contexts, and are often generically referred to as axionlike particles (ALPs). One particularly well-motivated example is the model-independent axion [25] in string theory. Since axions of this more general sort are, by and large, no less viable as dark-matter candidates than the QCD axion, it behooves us to extend our focus to encompass such fields as well.

For the remainder of this work, then, we will use the term "axion" to refer to any pseudoscalar field whose mass is generated by the instanton dynamics associated with an arbitrary non-Abelian gauge group G. This gauge group could be the SM $SU(3)$ color group, as it is for the QCD axion, but alternatively it could be some additional group which either resides in a hidden sector, or else confines at a very high scale. As with the QCD axion, any axion we consider will be assumed to be a pseudo-Nambu-Goldstone boson associated with the breaking of some global symmetry $U(1)_X$ at a scale f_X by the vacuum expectation value (VEV) of some scalar field. The axion field is assumed to couple to the field strength $\mathcal{G}_{\mu\nu}^a$ associated with G via a term of the form

$$\mathcal{L}_{\text{int}} \ni \frac{g_G^2 \xi}{32\pi f_X} a \mathcal{G}_{\mu\nu}^a \tilde{\mathcal{G}}^{\mu\nu a}, \qquad (2.9)$$

where a again denotes the axion field, g_G is the coupling constant associated with G, $\tilde{\mathcal{G}}^{\mu\nu a}$ is the dual of $\mathcal{G}^{\mu\nu a}$, and ξ is a model-dependent coefficient which parametrizes the strength of the effective interaction between the axion and the gauge fields. We will also assume that G goes through a confining phase transition at some scale Λ_G, and that a potential analogous to that appearing in Eq. (2.21) is thereby generated for a. In other words, this general axion couples to G in a manner completely analogous to that in which the QCD axion couples to the SM $SU(3)$. It therefore follows that all of the QCD-axion formalism outlined above continues to hold for axions in the broader sense of the word, provided one makes the substitutions $\Lambda_{\text{QCD}} \to \Lambda_G$, $g_3 \to g_G$, $f_{\text{PQ}} \to f_X$, etc., where appropriate.

There is, however, one crucial physical distinction between axions in general and the QCD axion in specific: for general axions, the confinement scale Λ_G is essentially a free parameter. The properties of such an axion are therefore far less constrained than those of a QCD axion, simply because the axion mass is not uniquely determined by f_X alone. Moreover, the vast majority of the experimental bounds on axions depend crucially on the charge assignments of the SM gauge fields under the global $U(1)_X$ symmetry. For a generic axion, these charges need not have any relationship to the $U(1)_{\text{PQ}}$ charge assignments for these fields. An important consequence of this is that a generalized axion need not couple directly to the gluon field at leading order. Moreover, other scenarios could be realized in such a framework that cannot arise for a QCD axion. For example, one can imagine a purely "photonic" axion which couples to the photon field at leading order, but not to the gluon field or to any of the SM fermions. In what follows, we will focus on several different concrete coupling scenarios. One of these will be such a photonic axion; another will be a "hadronic" axion which couples to the gluon and photon fields, but not directly to any of the SM fermions. However, we note that numerous other possibilities exist, and that the laboratory, astrophysical, and cosmological constraints on any given model depend sensitively on the couplings between the axion and the fields of the SM.

It is also worth noting that certain details of any scenario of this sort will depend on the details of the instanton dynamics associated with the particular gauge group G in question. The scaling behavior of $m_X(T)$ as a function of T, for example, may not be identical to the scaling behavior quoted in Eq. (2.6) for QCD instantons. However, none of these details plays a crucial role in the dark-matter phenomenology of the our model. We will therefore assume for the remainder of this work that, except for the values of Λ_G, g_G, etc., the standard axion results derived in the context of QCD-instanton dynamics apply to G-instanton dynamics as well.

B. Axions in extra dimensions

Having summarized the formalism applicable to a four-dimensional axion, we now consider how the situation changes when the axion in question is allowed to propagate in the extra-dimensional bulk of a theory with more than four dimensions. As was originally pointed out in Ref. [26], the dynamics of such an axion is far richer than that of a purely four-dimensional axion, due both to the presence of an entire KK tower of axion excitations and to a nontrivial mixing between these excitations due to the presence of brane-mass terms, which explicitly violate KK mode-number conservation. Indeed, as we shall see, it is those KK excitations which will constitute the dark-matter ensemble in our model, and it is their mixing which gives this ensemble the appropriate properties to be a viable dynamical dark-matter candidate. Of course scenarios involving large extra dimensions have many other attractive features as well: they provide a geometric interpretation of the hierarchy between the weak scale and the Planck scale [27–29], between the weak scale and the grand-unification scale [8], and between the weak scale and the string scale [8,30]. Moreover, a higher-dimensional axion field can be accommodated quite naturally in such a brane/bulk framework. Indeed, while only gravity is required to propagate in the bulk, the propagation of SM-gauge-singlet fields there, including axions of all varieties, is, in a sense, almost expected.

In what follows, we present the setup for a generic axion field in the bulk. This parallels the setup for a QCD axion put forth in Ref. [26]. For concreteness, we choose to focus on the case in which the axion is allowed to propagate in a single, large extra dimension compactified on a $S_1/\mathbb{Z}_2$ orbifold of radius R, while the fields of the SM and the gauge fields associated with the additional symmetry group G are confined to the brane located at $x_5 = 0$. However, we emphasize that the setup described here can easily be extended to scenarios in which the axion in question is allowed to propagate in multiple extra dimensions, or in which the background geometry is more complicated [31].

At scales below the weak scale but above the confinement scale Λ_G, the effective action for a bulk axion in five dimensions takes the form

$$S_{\text{eff}} = \int d^4x \int_0^{2\pi R} dx_5 \left[\frac{1}{2} \partial_M a \partial^M a + \delta(x_5)(\mathcal{L}_{\text{brane}} + \mathcal{L}_{\text{int}}) \right].$$
(2.10)

Here, we have divided the brane-localized terms in the Lagrangian into two parts. The first, $\mathcal{L}_{\text{brane}}$, contains the terms involving the brane fields alone—both the fields of the SM and any additional fields, including the gauge fields associated with the gauge group G. The second, $\mathcal{L}_{\text{int}}$, contains the interaction terms involving the brane-localized fields and the five-dimensional axion. This second piece is given by

$$\mathcal{L}_{\text{int}} = \frac{g_G^2 \xi}{32\pi^2 f_X^{3/2}} a G_{\mu\nu}^a \tilde{G}^{a\mu\nu} + \sum_i \frac{c_i}{f_X^{3/2}} (\partial_\mu a) \bar{\psi}_i \gamma^\mu \gamma^5 \psi_i$$
$$+ \frac{g_s^2 c_g}{32\pi^2 f_X^{3/2}} a G_{\mu\nu}^a \tilde{G}^{a\mu\nu} + \frac{e^2 c_\gamma}{32\pi^2 f_X^{3/2}} a F_{\mu\nu} \tilde{F}^{\mu\nu},$$
(2.11)

where e and g_s are the respective couplings for $U(1)_{\text{EM}}$ and $SU(3)$ color, and f_X is the fundamental five-dimensional scale associated with the breaking of $U(1)_X$ (the analogue of the Peccei-Quinn scale f_{PQ} in Ref. [26]).

The first term in $\mathcal{L}_{\text{int}}$ is the requisite coupling between the five-dimensional axion a and the gauge fields of G. The second term represents the derivative couplings between the five-dimensional axion a and the SM fermion fields ψ_i, with model-dependent coefficients c_i that depend on the $U(1)_X$ charges of the ψ_i. The remaining two terms represent the interactions between the axion and the gluon and photon fields, the field-strength tensors for which are here, respectively, denoted $G^{\mu\nu a}$ and $F^{\mu\nu}$, with (once again model-dependent) coefficients c_γ and c_g.

The five-dimensional axion field can be represented as a tower of KK excitations via the decomposition

$$a(x^\mu, x_5) = \frac{1}{\sqrt{2\pi R}} \sum_{n=0}^{\infty} r_n a_n(x^\mu) \cos\left(\frac{nx_5}{R}\right),$$
(2.12)

where the factor

$$r_n \equiv \begin{cases} 1 & \text{for } n = 0 \\ \sqrt{2} & \text{otherwise} \end{cases}$$
(2.13)

ensures that the kinetic term for each mode is canonically normalized. Substituting this expression into Eq. (2.11) and integrating over x_5, we obtain

$$S_{\text{eff}} = \int d^4x \left[\sum_{n=0}^{\infty} \left(\frac{1}{2} \partial_\mu a_n \partial^\mu a_n + \frac{g_G^2 \xi}{32\pi^2 \hat{f}_X} r_n a_n G_{\mu\nu}^a \tilde{G}^{a\mu\nu} \right. \right.$$
$$+ \sum_i \frac{c_i}{\hat{f}_X} r_n (\partial_\mu a_n) \bar{\psi}_i \gamma^\mu \gamma^5 \psi_i + \frac{g_s^2 c_g}{32\pi^2 \hat{f}_X} r_n a_n G_{\mu\nu}^a \tilde{G}^{a\mu\nu}$$
$$\left. \left. + \frac{e^2 c_\gamma}{32\pi^2 \hat{f}_X} r_n a_n F_{\mu\nu} \tilde{F}^{\mu\nu} \right) - V(a) \right],$$
(2.14)

where

$$V(a) = \sum_{n=0}^{\infty} \frac{1}{2} \frac{n^2}{R^2} a_n^2,$$
(2.15)

and where the quantity $\hat{f}_X$, defined by the relation

$$\hat{f}_X^2 \equiv 2\pi R f_X^3,$$
(2.16)

represents the effective four-dimensional $U(1)_X$-breaking scale. Note that each mode in the KK tower couples to the SM fields with a strength inversely proportional to $\hat{f}_X$. Note also that at these scales, the axion mass-squared matrix

$$\mathcal{M}_{mn}^2 \equiv \frac{\partial^2 V(a)}{\partial a_m \partial a_n}$$
(2.17)

is purely diagonal.

The effective action in Eq. (2.14) is valid at high scales where $T \gg \Lambda_G$. Around $T \sim \Lambda_G$, however, instanton effects give rise to an additional contribution to the effective axion potential. In the low-temperature regime, the full potential takes the form

$$V(a) = \sum_{n=0}^{\infty} \frac{1}{2} \frac{n^2}{R^2} a_n^2$$
$$+ \frac{g_G^2}{32\pi^2} \Lambda_G^4 \left[1 - \cos\left(\frac{\xi}{\hat{f}_X} \sum_{n=0}^{\infty} r_n a_n + \bar{\Theta}_G \right) \right],$$
(2.18)

where $\bar{\Theta}_G$ is the analogue of the QCD theta-parameter $\bar{\Theta}$. Minimizing the potential yields the vacuum configuration $\langle a_0 \rangle = \hat{f}_X(-\bar{\Theta}_G + \pi\ell)/\xi$ for $\ell \in 2\mathbb{Z}$, with $\langle a_n \rangle = 0$ for all $n > 0$. This additional potential term modifies the axion mass-squared matrix at scales $T \lesssim \Lambda_G$ to

$$\mathcal{M}_{mn}^2 = M_c^2 n^2 \delta_{mn}$$
$$+ \frac{g_G^2 \xi^2}{32\pi^2} \frac{\Lambda_G^4}{\hat{f}_X^2} r_m r_n \cos\left(\frac{\xi}{\hat{f}_X} \sum_{k=0}^{\infty} r_k a_k + \bar{\Theta}_G \right),$$
(2.19)

where $M_c \equiv 1/R$ is the compactification scale. We see here that the terms originating from the instanton-induced potential in Eq. (2.18) include off-diagonal contributions, which result in mixings among the KK eigenstates. In the vicinity of the minimum of $V(a_n)$, the axion mass-squared matrix above takes the form [26]

$$\mathcal{M}^2 = m_X^2 \begin{pmatrix} 1 & \sqrt{2} & \sqrt{2} & \sqrt{2} & \dots \\ \sqrt{2} & 2+y^2 & 2 & 2 & \dots \\ \sqrt{2} & 2 & 2+4y^2 & 2 & \dots \\ \sqrt{2} & 2 & 2 & 2+9y^2 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix},$$

(2.20)

where

$$y \equiv \frac{M_c}{m_X} \quad \text{and} \quad m_X^2 \equiv \frac{g_G^2 \xi^2}{32\pi^2} \frac{\Lambda_G^4}{\hat{f}_X^2}.$$

(2.21)

The eigenvalues λ^2 of this mass-squared matrix are the set of solutions to the transcendental equation

$$\frac{\pi \lambda m_X}{y} \cot\left(\frac{\pi \lambda}{m_X y}\right) = \lambda^2.$$

(2.22)

The normalized mass eigenstate a_λ corresponding to each λ may be written as a sum of the KK eigenstates a_n:

$$a_\lambda = \sum_{n=0}^{\infty} U_{\lambda n} a_n \equiv \sum_{n=0}^{\infty} \left(\frac{r_n \tilde{\lambda}^2}{\tilde{\lambda}^2 - n^2 y^2}\right) A_\lambda a_n,$$

(2.23)

where $\tilde{\lambda} \equiv \lambda/m_X$, and where

$$A_\lambda \equiv \frac{\sqrt{2}}{\tilde{\lambda}} [1 + \tilde{\lambda}^2 + \pi^2/y^2]^{-1/2}.$$

(2.24)

The quantity A_λ can be shown to obey the sum rules [26]

$$\sum_\lambda A_\lambda^2 = 1, \qquad \sum_\lambda \tilde{\lambda}^2 A_\lambda^2 = 1,$$

(2.25)

which follow directly from the unitarity of $U_{\lambda n}$. Upon rewriting Eq. (2.14) in this mass eigenbasis, we obtain the axion effective action at temperatures $T \lesssim \Lambda_G$, which, up to $\mathcal{O}(a_\lambda^6/\hat{f}_X^6)$, is given by

$$S_{\text{eff}} = \int d^4x \Bigg[\sum_\lambda \Big(\frac{1}{2} \partial_\mu a_\lambda \partial^\mu a_\lambda - \frac{1}{2} \tilde{\lambda}^2 m_X^2 a_\lambda^2 + \frac{e^2 c_\gamma \tilde{\lambda}^2 A_\lambda}{32\pi^2 \hat{f}_X} a_\lambda F_{\mu\nu} \tilde{F}^{\mu\nu} + \frac{g_s^2 c_g \tilde{\lambda}^2 A_\lambda}{32\pi^2 \hat{f}_X} a_\lambda G_{\mu\nu}^a \tilde{G}^{\mu\nu a}$$
$$+ \sum_i \frac{c_i \tilde{\lambda}^2 A_\lambda}{\hat{f}_X} (\partial_\mu a_\lambda) \bar{\psi}_i \gamma^\mu \gamma^5 \psi_i \Big) + \frac{g_G^2 \xi^4 \Lambda_G^4}{768\pi^2 \hat{f}_X^4} \sum_{\lambda_i, \lambda_j, \lambda_k, \lambda_\ell} \tilde{\lambda}_i^2 \tilde{\lambda}_j^2 \tilde{\lambda}_k^2 \tilde{\lambda}_\ell^2 A_{\lambda_i} A_{\lambda_j} A_{\lambda_k} A_{\lambda_\ell} a_{\lambda_i} a_{\lambda_j} a_{\lambda_k} a_{\lambda_\ell} \Bigg].$$

(2.26)

The quartic axion self-interaction terms shown above originate from the instanton-induced potential in Eq. (2.18). Other, higher-order terms not shown may be safely neglected when $T \ll \hat{f}_X$.

If a nontrivial coupling exists between the bulk axion and the gluon field, effective interactions will also arise between the a_λ and the hadron fields at temperatures below Λ_{QCD}. (In the case of the QCD axion, of course, such couplings are mandatory.) The Lagrangian which describes these interactions is just the five-dimensional analogue of Eq. (2.6):

$$\mathcal{L}_{\text{had}} = \tilde{\lambda}^2 A_\lambda \frac{C_{a\pi}}{f_\pi \hat{f}_X} (\partial_\mu a_\lambda) [(\partial^\mu \pi^+)\pi^- \pi^0 + (\partial^\mu \pi^-)\pi^+ \pi^0 - 2(\partial^\mu \pi^0)\pi^+ \pi^-] + \tilde{\lambda}^2 A_\lambda \frac{C_{an}}{\hat{f}_X} (\partial_\mu a_\lambda) \bar{n} \gamma^\mu \gamma^5 n$$
$$+ \tilde{\lambda}^2 A_\lambda \frac{C_{ap}}{\hat{f}_X} (\partial_\mu a_\lambda) \bar{p} \gamma^\mu \gamma^5 p + i\tilde{\lambda}^2 A_\lambda \frac{C_{a\pi N}}{f_\pi \hat{f}_X} (\partial_\mu a_\lambda) [\pi^+ \bar{p} \gamma^\mu n - \pi^- \bar{n} \gamma^\mu p],$$

(2.27)

where the coefficients $C_{a\pi}$, C_{an}, etc., depend on the details of the theory, and may differ from those given in Eqs. (2.7) and (2.8).

As discussed in Ref. [13], the mass spectrum of the model reduces to a reasonably simple form in certain limiting cases which depend on the value of the ratio y defined in Eq. (2.21). The first of these is weakly-mixed regime, in which $y \gg 1$. In this regime, the a_λ are all very nearly equivalent to the KK modes a_n, with masses $\lambda \approx nM_c$, where n is an integer. The extent to which any given value of λ differs from M_c is set by the size of the off-diagonal terms in Eq. (2.20), and, in particular, the lightest

mass eigenstate a_{λ_0} has a mass $\lambda_0 \approx M_c/y = m_X$. In short, the KK tower essentially comprises a single light mode plus a tower of massive KK excitations of that mode. In the extreme limit, in which $M_c \rightarrow \infty$, the theory reduces to an effectively four-dimensional theory with a single light axion whose mass is precisely equal to m_X, as expected.

In the opposite limit, in which $y \ll 1$, the situation is markedly different [26]. The lighter mass eigenstates in the theory have masses $\lambda \approx (n + 1/2)M_c \ll m_X$ (where $n = \{0, 1, 2, \ldots\}$ is an integer), while the heavier mass eigenstates have masses $\lambda \approx nM_c$. The transition region between the two regimes occurs at around

$$\lambda \sim \lambda_{\text{trans}} \equiv \frac{\pi m_X^2}{M_c}, \qquad (2.28)$$

which corresponds to a value $n = \pi/y^2$. In this regime, the states with masses below this threshold are highly mixed, owing to the large, off-diagonal terms in Eq. (2.20) proportional to m_X. We dub this the "strongly-mixed" regime. In this latter regime, as we shall soon see, the full KK tower plays a far larger role in the dark-matter phenomenology of a given model than in weakly-mixed scenarios, in which the dark-matter phenomenology is, more or less, the phenomenology of the zero mode. Indeed, the bulk-axion scenarios which give rise to dynamical dark-matter ensembles in which the full tower contributes significantly tend to be those in which y is small.

It is evident from Eqs. (2.26) and (2.27) that the parameter combination $\tilde{\lambda}^2 A_\lambda / \hat{f}_X$ plays a critical role in bulk-axion dynamics. Indeed, it is this combination which determines the strength of the interaction between a given mass eigenstate a_λ and any of the SM fields. It turns out to be phenomenologically quite significant that not all a_λ couple to the fields of the SM with the same strength. We discuss the impact of such a coupling structure on the decay properties of the tower states in Sec. III, and summarize its impact on phenomenological constraints in Sec. VI (a more detailed analysis of which can be found in Ref. [32]). In addition, a plot of how A_λ and $\tilde{\lambda}^2 A_\lambda$ depend on λ is provided in Ref. [13]. It is worth remarking that the nonuniversality of the a_λ couplings is yet another direct consequence of the nontrivial mixing between axion KK modes implied by Eq. (2.20). This effect does not arise for bulk fields in the absence of such mixing: the couplings of the KK excitations of the graviton to the SM fields in theories of this sort, for example, are identical for all modes.

In summary, our model for dynamical dark matter consists of an axion propagating in the bulk of an extra dimension of radius R, with the SM living on a brane. From the perspective of an observer on the brane, our dynamical dark-matter ensemble consists of the KK modes of this bulk-axion field. As we have discussed above, our model involves three important dimensionful parameters: $\hat{f}_X$, M_c, and Λ_G. We have also shown that the physics of this model depends crucially on one particular dimensionless combination of these parameters, namely y, which governs the extent to which the individual KK modes mix when forming the constituents of our dynamical dark-matter ensemble.

III. CHARACTERIZING THE CONSTITUENTS: DECAY WIDTHS

Now that we have reviewed the setup underlying our model for dynamical dark matter, we may begin to assess its phenomenological ramifications. As discussed in Ref. [13], the essence of the dynamical dark-matter framework lies in the balance between the decay widths and relic abundances of the fields which contribute to Ω_{CDM}. Therefore, our principal aim must be to evaluate the decay widths Γ_λ and relic abundances Ω_λ of the fields a_λ which our dark-matter ensemble comprises, and examine how these two quantities scale with λ. In this section, we focus on decays: we calculate the partial-width contribution associated with each of the potentially relevant decay channels for a generic a_λ and assess how the total width Γ_λ depends on the dimensionful parameters of the model, namely $\hat{f}_X$, M_c, and Λ_G. In the subsequent section, we focus on abundances.

A. Decays to standard-model states

We begin our discussion of axion decays by computing the decay widths of the a_λ directly to SM states on the brane. The first step is to derive Feynman rules for the relevant interactions, which can be obtained directly from the interaction terms given in Eq. (2.26) and (2.27). For those a_λ with λ below a few GeV, the relevant vertices are:

Here, and throughout the rest of this work, the symbol M_P represents the effective, four-dimensional *reduced* Planck mass. Since the coupling of each a_λ to the SM fields is suppressed by $\hat{f}_X$, the contribution to the total decay width Γ_λ of each a_λ which comes from decays to SM fields will be suppressed by $\hat{f}_X^{-2}$. At low temperatures, the most relevant decay processes (depending, of course, on the precise values of c_i and c_γ in any given model) will be $a_\lambda \rightarrow \gamma\gamma$, $a_\lambda \rightarrow e^+ e^-$, and $a_\lambda \rightarrow \nu_i \bar{\nu}_i$, where $i = \{1, 2, 3\}$ labels the three light neutrino mass eigenstates.

The one decay channel which is kinematically accessible for all axion modes, regardless of their mass (provided c_γ is nonvanishing), is $a \rightarrow \gamma\gamma$. The partial decay width of an axion mass eigenstate a_λ into a pair of photons is

$$\Gamma(a_\lambda \to \gamma\gamma) = \frac{c_\gamma^2 \alpha^2 \lambda^3}{256\pi^3 \hat{f}_X^2}(\tilde{\lambda}^2 A_\lambda)^2 = G_\gamma (\tilde{\lambda}^2 A_\lambda)^2 \frac{\lambda^3}{\hat{f}_X^2}, \qquad (3.1)$$

where $\alpha \equiv e^2/4\pi$, and where we have defined the quantity $G_\gamma \equiv c_\gamma^2 \alpha^2/256\pi^3$. Note that $\Gamma(a_\lambda \to \gamma\gamma)$ includes an overall factor $(\tilde{\lambda}^2 A_\lambda)^2$. This factor is a direct consequence of mixing in the axion mass matrix, and it appears universally in all partial-width expressions for axion decays to states on the brane. For those modes with masses $\lambda \gtrsim \pi m_X^2/M_c$, this factor is $\mathcal{O}(1)$; however, for those modes with $\lambda \lesssim \pi m_X^2/M_c$, this factor can be much smaller. We therefore see that the mixing factor $(\tilde{\lambda}^2 A_\lambda)^2$ suppresses the decay widths of the lighter a_λ, while leaving the widths of the heavier a_λ unsuppressed. This decay-width suppression for the light modes plays a crucial role in bulk-brane models of dynamical dark matter, as discussed in Ref. [13]: since these light modes also turn out to have the largest relic abundances, their decays are generally the most dangerous from a phenomenological perspective.

As λ increases, additional decay channels open up in which a given a_λ decays to a light fermion-antifermion pair, provided that direct couplings exist between the axion and the fermionic species ψ_i in question. The partial width for any decay of this sort is given by

$$\Gamma(a_\lambda \to \psi_i \bar{\psi}_i) = \frac{c_i^2 \lambda m_\psi^2}{2\pi \hat{f}_X^2}(\tilde{\lambda}^2 A_\lambda)^2 \left(1 - \frac{4m_\psi^2}{\lambda^2}\right)^{1/2}, \qquad (3.2)$$

where m_ψ is the mass of the fermion in question. Note that if the five-dimensional axion field couples to the SM neutrinos, the Γ_λ will have a nontrivial dependence on the neutrino mass spectrum. While the precise masses of the three neutrino species are as yet unknown, measurements of the solar and atmospheric squared-mass splittings place lower limits on two of the three m_{ν_i}. The current best-fit values for these splittings are [33]

$$\Delta m_\odot^2 = 7.59^{+0.19}_{-0.21} \times 10^{-5} \text{ eV}^2,$$
$$|\Delta m_A^2| = 2.43^{+0.13}_{-0.13} \times 10^{-3} \text{ eV}^2. \qquad (3.3)$$

In what follows we shall assume a normal hierarchy, and we will take the mass of the lightest neutrino to be vanishingly small. In this case, the masses of the heavier two neutrinos are $m_{\nu_2} \simeq 8.7 \times 10^{-3}$ eV and $m_{\nu_3} \simeq 4.9 \times 10^{-2}$ eV, which are comparable to the lower bound [34]

$$M_c \gtrsim 3.9 \times 10^{-3} \text{ eV} \qquad (3.4)$$

on the compactification scale M_c from modified-gravity experiments. Therefore, if M_c lies only slightly above this bound, the masses λ of certain light a_λ will be comparable to $m_{\nu_{2,3}}$. The partial width for the decays of those a_λ to neutrinos can therefore in principle be quite large compared to their partial widths for decays to photons, as can be seen by comparing Eqs. (3.1) and (3.2). Note that if $c_{\nu_i} = 0$ for all i, the decay-width loses all sensitivity to the neutrino mass spectrum. This is indeed the case for the photonic and

hadronic axions which will serve as our primary examples in what follows.

As λ increases still further, decays of the a_λ to hadrons become kinematically accessible—provided, of course, that either $c_g \neq 0$, or else that $c_{q_i} \neq 0$ for some quark species q_i. The lowest such threshold is that for decays of the form $a \to \pi^+ \pi^- \pi^0$, which is kinematically allowed whenever $\lambda > 2m_{\pi^\pm} + m_\pi^0$. The relevant interaction vertex is that appearing in the top line of Eq. (2.27), and the corresponding contribution to the decay width of a_λ from this three-body decay takes the form

$$\Gamma(a_\lambda \to \pi^+ \pi^- \pi^0) = \frac{C_{a\pi}^2 (\tilde{\lambda}^2 A_\lambda)^2}{1024\pi^3 \lambda^3 f_\pi^2 \hat{f}_X^2} I(\lambda), \qquad (3.5)$$

where $I(\lambda)$ denotes the phase-space integral

$$I(\lambda) = \int_{4m_{\pi^\pm}^2}^{(\lambda - m_{\pi_0})^2} dm_{12}^2 (\lambda^2 + m_{12}^2 - m_{\pi_0}^2)^2 \left(1 - \frac{4m_{\pi^\pm}^2}{m_{12}^2}\right)^{1/2}$$
$$\times [m_{12}^4 + 2m_{12}^2(\lambda^2 - 3m_{\pi^0}^2) + (\lambda^2 - m_{\pi^0}^2)^2]^{1/2}, \qquad (3.6)$$

with integration variable $m_{12}^2 \equiv (k_+ + k_-)^2$. For $\lambda \gg 2m_{\pi^\pm} + m_{\pi^0}$, this expression takes the asymptotic form

$$\Gamma(a_\lambda \to \pi^+ \pi^- \pi^0) \approx (2.07 \times 10^{-2}) C_{a\pi}^2 \left(\frac{\lambda^5}{f_\pi^2 \hat{f}_X^2}\right)(\tilde{\lambda}^2 A_\lambda)^2. \qquad (3.7)$$

In practice, this asymptotic expression is a good approximation for $\Gamma(a_\lambda \to \pi^+ \pi^- \pi^0)$ as long as λ roughly exceeds a few GeV.

For even larger values of λ, decays to nucleons become kinematically accessible. In the present treatment, however, any a_λ with masses this large will not play a significant role in the phenomenology of such "dark-tower" scenarios, nor will they have a significant impact on the observational and experimental constraints on such scenarios. Indeed, as was observed in Ref. [35] for the related case of KK-graviton decays, such modes are innocuous precisely *because* of the large contributions to their decay widths from hadronic decays. The axion case under consideration here differs qualitatively from this KK-graviton case only in that the decay width Γ_λ of each a_λ contains an additional factor $(\tilde{\lambda}^2 A_\lambda)^2$. However, as we shall see in Sec. VI, the quantity $\pi m_X^2/M_c$ is never much larger than a few GeV in realistic axion models of dynamical dark matter. Consequently, this factor will be $\mathcal{O}(1)$ for modes with λ roughly exceeding a few GeV—those modes for which decays to nucleons are kinematically allowed—and the partial widths for those decays will therefore be unsuppressed. We therefore refrain from explicitly calculating the partial-width contribution from such decays, and simply acknowledge that any a_λ with λ roughly exceeding a few GeV will decay quite early—i.e., before the big-bang-nucleosynthesis (BBN) epoch.

Thus far, we have obtained partial-width expressions for all of the relevant decay channels for those a_λ with λ less than roughly a few GeV directly into final states involving SM fields alone. The total width for the decays of a_λ into this set of final states is obtained by combining the expressions in Eqs. (3.1), (3.2), and (3.5):

$$\Gamma_\lambda = \frac{\lambda^3}{8\pi \hat{f}_X^2}(\tilde{\lambda}^2 A_\lambda)^2\left[\frac{\alpha^2 c_\gamma^2}{32\pi^2} + \sum_i \Theta(\lambda - 2m_{\psi_i})\frac{4c_i^2 m_{\psi_i}^2}{\lambda^2}\left(1 - \frac{4m_{\psi_i}^2}{\lambda^2}\right)^{1/2} + \Theta(\lambda - 2m_{\pi^\pm} - m_{\pi^0})\frac{C_{a\pi}^2 I(\lambda)}{128\pi^2 f_\pi^2 \lambda^6}\right], \quad (3.8)$$

where the Heaviside functions $\Theta(\lambda - 2m_\psi)$ and $\Theta(\lambda - 2m_{\pi^\pm} - m_{\pi^0})$ enforce that only kinematically accessible decays contribute in the sum.

B. Intraensemble decays

Up to this point, we have only been considering the decay-width contributions from decays of a given a_λ directly to the fields of the SM. We have yet to address the issue of *intraensemble* decays—i.e., decays in which a given state in the dynamical dark-matter ensemble decays to a final state containing one or more other, lighter states in that ensemble. Indeed, as discussed in Ref. [13], such decays can have a significant impact on the phenomenology of a given dynamical dark-matter model: not only will they alter the individual relic abundances Ω_i of the particles ϕ_i in a given ensemble, but they will also alter the phase-space distributions $f_i(\vec{p}_i, t)$ of those particles, potentially generating a sizable population of ϕ_i with relativistic three-momenta $\vec{p}_i$. For these reasons, it is crucial to assess whether such decays occur at a substantial rate in the bulk-axion model presented here, or whether the net rate for these decays is negligible, in which case the quantity Γ_λ given in Eq. (3.8) truly represents the total decay width of a given a_λ with λ less than a few GeV.

In the model under discussion here, the dark sector properly comprises KK axions, KK gravitons, and graviscalars. A complete description of the dynamics of the ensemble would therefore involve solving the coupled system of Boltzmann equations which describes the evolution of the respective phase-space distributions $f_\lambda(\vec{p}_\lambda, t)$, $f_n(\vec{p}_n, t)$, and $f_s(\vec{p}_s, t)$ for the various axion modes a_λ, KK gravitons $G_{\mu\nu}^{(n)}$, and graviscalars φ_s, as discussed in the Appendix of Ref. [13]. In the present work, our aim will not be to solve the Boltzmann equations in complete generality, but rather to demonstrate that the effects of intraensemble decays on the abundances and phase-space distributions of the ϕ_i are sufficiently small that they may be safely neglected for any otherwise phenomenologically reasonable choice of model parameters. Some of these effects—for example, the depletion of Ω_λ for a given a_λ due to intraensemble decays—depend only on the net contribution $\Gamma_\lambda^{(\mathrm{IE})}$ to the width of a given a_λ, obtained by summing over the partial widths from all such decays. On the other hand, certain other effects, such as the increase in the abundances Ω_λ and alteration of the phase-space distributions $f_\lambda(\vec{p}_\lambda, t)$ of the lighter a_λ due to the decays

of the heavier a_λ, depend on these partial widths in a different manner. Determining the precise magnitude of these effects therefore requires a more thorough analysis of the Boltzmann equations. Nevertheless, $\Gamma_\lambda^{(\mathrm{IE})}$ can still be useful as a rough rubric for assessing whether or not the effects in question are likely to be significant. In this work, then, we simply demonstrate that $\Gamma_\lambda^{(\mathrm{IE})}$ for any given a_λ is sufficiently small in comparison with the result in Eq. (3.8) in otherwise phenomenologically reasonable regions of model parameter space that its effect on total decay widths may safely be neglected. This result provides a good initial indication that the additional effects mentioned above are also unimportant. A more rigorous justification for neglecting these effects will appear in Ref. [32].

We begin our discussion of intraensemble decays by discussing the partial-width contribution arising from the decays of a given a_λ into multiple, lighter axion modes. The leading contribution to this partial width comes from the quartic interaction term appearing in Eq. (2.26), which gives rise to decays of the form $a_\lambda \rightarrow a_{\lambda_1} a_{\lambda_2} a_{\lambda_3}$, where a_{λ_i}, with $i = \{1, 2, 3\}$, are lighter axion states whose masses satisfy the constraint $\lambda \gtrsim \lambda_1 + \lambda_2 + \lambda_3$. The Feynman rule for the corresponding four-point interaction vertex is
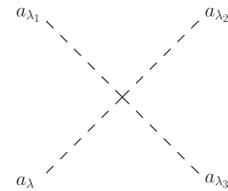
from which the contribution to Γ_λ from this three-body decay is found to be

$$\Gamma(a_\lambda \rightarrow a_{\lambda_1} a_{\lambda_2} a_{\lambda_3}) = \frac{4\pi}{\lambda^3}\left(\frac{g_G \xi^2 \Lambda_G^2}{32\pi^2 \hat{f}_X^2}\right)^4 (\tilde{\lambda}^2 A_\lambda)^2 \left(\prod_{i=1}^{3} \tilde{\lambda}_i^2 A_{\lambda_i}\right)^2$$
$$\times \int_{(m_{12}^2)_{\min}}^{(m_{12}^2)_{\max}} \int_{(m_{23}^2)_{\min}}^{(m_{23}^2)_{\max}} dm_{12}^2 dm_{23}^2,$$

$$(3.9)$$

where now $m_{12}^2 = (p_1 + p_2)^2$ and $m_{23}^2 = (p_2 + p_3)^2$, with p_i representing the four-momentum of the final-state a_{λ_i}. The limits of integration for the dm_{12} integral are $(m_{12})_{\max}^2 = (\lambda - \lambda_3)^2$ and $(m_{12})_{\min}^2 = (\lambda_1 + \lambda_2)^2$, and

since we are primarily interested in the parametric dependence of the partial width on λ, M_c, etc., it will be sufficient for our purposes to construct an upper bound on $\Gamma(a_\lambda \rightarrow a_{\lambda_1} a_{\lambda_2} a_{\lambda_3})$ by setting $(m_{23})^2_{\text{max}} \rightarrow (\lambda - \lambda_1)^2$ and $(m_{23})^2_{\text{min}} \rightarrow (\lambda_2 + \lambda_3)^2$. Doing so, we obtain the result

$$
\begin{aligned}
&\Gamma(a_\lambda \rightarrow a_{\lambda_1} a_{\lambda_2} a_{\lambda_3}) \\
&\leq \frac{4\pi}{\lambda^3}\left(\frac{g_G \xi^2 \Lambda_G^2}{32\pi^2 \hat{f}_X^2}\right)^4 (\tilde{\lambda}^2 A_\lambda)^2 \left(\prod_{i=1}^{3} \tilde{\lambda}_i^2 A_{\lambda_i}\right)^2 \\
&\quad \times (\lambda_1 + \lambda_2 + \lambda_3 - \lambda)^2 \\
&\quad \times (\lambda^2 - \lambda_1^2 + \lambda_2^2 - \lambda_3^2 + 2\lambda\lambda_2 + 2\lambda_1\lambda_3). \quad (3.10)
\end{aligned}
$$

Note that the asymmetry of this expression under permutations of the λ_i is due to the asymmetric limits of integration we have adopted in order to construct this bound. In order to obtain the total contribution $\Gamma_\lambda(a \rightarrow 3a)$ to the partial width of a given a_λ from decays of this form, we need to sum over the contributions from all kinematically allowed decays of the form $a_\lambda \rightarrow a_{\lambda_1} a_{\lambda_2} a_{\lambda_3}$. Next, we approximate the sums over the different allowed final-state axions with integrals over $d\lambda_1$, $d\lambda_2$, and $d\lambda_3$. Furthermore, Eq. (2.24) implies that $\tilde{\lambda}^2 A_\lambda < \sqrt{2}$ for all λ. Therefore, in order to obtain an upper bound on $\Gamma_\lambda(a \rightarrow 3a)$, we make the replacements $\tilde{\lambda}_i^2 A_{\lambda_i} \rightarrow \sqrt{2}$ and $\tilde{\lambda}^2 A_\lambda \rightarrow \sqrt{2}$. Doing so, we obtain our final result:

$$
\Gamma_\lambda(a \rightarrow 3a) \leq \frac{g_G^4 \xi^8}{45(4\pi)^7} \frac{\lambda^4}{M_c^3}\left(\frac{\Lambda_G}{\hat{f}_X}\right)^8. \quad (3.11)
$$

Given this result, we are now prepared to address the question of whether axion decays to other bulk axions can ever contribute significantly to the total width Γ_λ of any a_λ. For example, the bound in Eq. (3.11) implies that the ratio of $\Gamma_\lambda(a \rightarrow 3a)$ to the decay rate $\Gamma_\lambda(a \rightarrow \gamma\gamma)$ given in Eq. (3.1) is bounded from above by

$$
\begin{aligned}
\frac{\Gamma_\lambda(a \rightarrow 3a)}{\Gamma_\lambda(a \rightarrow \gamma\gamma)} &\leq \frac{4g_G^4 \xi^8}{45(4\pi)^4 \alpha^2 c_\gamma^2}\left(\frac{\lambda \Lambda_G^8}{M_c^3 \hat{f}_X^6}\right) \\
&\approx (6.69 \times 10^{-2}) \times (g_G \xi^2)^4 \left(\frac{\lambda \Lambda_G^8}{M_c^3 \hat{f}_X^6}\right). \quad (3.12)
\end{aligned}
$$

As we shall see in Sec. VI, for $\mathcal{O}(1)$ values of ξ and g_G, the phenomenologically preferred ranges for $\hat{f}_X$ and Λ_G turn out to be $\hat{f}_X \sim 10^{14}$–10^{15} GeV and $\Lambda_G \sim 10^3$–10^5 GeV, while M_c is bounded from below by Eq. (3.4). Within this parameter-space regime, we find that the ratio in Eq. (3.12) will be vanishingly small, as desired, unless $\lambda \gtrsim 10^{10}$ GeV. Since this is far larger than the cutoff scale f_X in this same regime, we conclude that decays of the form $a_\lambda \rightarrow a_{\lambda_1} a_{\lambda_2} a_{\lambda_3}$ will not play a significant role in the phenomenology of realistic bulk axion models of dynamical dark matter. Thus such decays can be safely neglected in computing the total decay width of a given a_λ.

Not only can the a_λ decay to final states comprising lighter axion modes alone, but they can also decay into final states which include other bulk states. In the minimal bulk-axion theory under discussion here, these include KK graviscalars and KK gravitons. (Note that the vector degrees of freedom $h_{\mu 5}^{(n)}$ with $n > 0$ in the gravity multiplet do not couple to the a_λ in the linearized-gravity limit in the unitary gauge, and the zero-mode $h_{\mu 5}^{(0)}$ vanishes due to the orbifold projection.) Therefore, we must also assess whether decay channels involving these KK gravitons and KK graviscalars can provide an appreciable contribution to $\Gamma_\lambda^{(\text{IE})}$. In the five-dimensional theory under discussion here, in the unitary gauge, the only physical graviscalar present is a single radion mode, which we assume here to be sufficiently massive (e.g., as the result of some stabilization mechanism) as not to be relevant for a_λ decays. As for decays involving KK gravitons in the final state, a rough upper bound on their contribution to Γ_λ for the case of a single, flat extra dimension will be given in Ref. [32] within the framework of linearized gravity. As we will see, the leading contribution comes from two-body decays of the form $a_\lambda \rightarrow G_{\mu\nu}^{(n)} a_{\lambda'}$, where $G_{\mu\nu}^{(n)}$ denotes the KK graviton with KK mode number n. The total contribution from decays of this sort, summed over all kinematically accessible combinations of n and λ', is found to be approximately [32]

$$
\begin{aligned}
\Gamma(a_\lambda \rightarrow Ga) &\lesssim \frac{8m_X^4(\tilde{\lambda}^2 A_\lambda)^2}{9\pi\lambda^3 M_c^2 M_P^2}\int_0^\lambda d\lambda'(\tilde{\lambda}'^2 A_{\lambda'})^2(\lambda + \lambda') \\
&\quad \times [(\lambda^2 + \lambda'^2)E(x_\lambda) - 2\lambda\lambda'K(x_\lambda)], \quad (3.13)
\end{aligned}
$$

where $K(x_\lambda)$ and $E(x_\lambda)$ respectively denote the complete elliptic integrals of the first and second kind, with $x_\lambda \equiv (\lambda - \lambda')^2/(\lambda + \lambda')^2$.

In order to compare $\Gamma(a_\lambda \rightarrow Ga)$ to the rate for a_λ decays to SM fields, it is necessary to integrate Eq. (3.13) numerically, as a function of $\hat{f}_X$, Λ_G, and M_c. The results of such an analysis are given in Ref. [32]. Here, however, to illustrate our point, we simply choose a set of benchmark values typical of a phenomenologically consistent scenario for which $\Gamma(a_\lambda \rightarrow Ga)/\Gamma(a_\lambda \rightarrow \gamma\gamma)$ is roughly maximal. Specifically, we take $\Lambda_G = 1$ TeV and $M_c = 10^{-11}$ GeV, with $g_G = \xi = 1$. We then find that $\Gamma(a_\lambda \rightarrow Ga)$ always remains several orders of magnitude smaller than $\Gamma(a_\lambda \rightarrow \gamma\gamma)$ for all values of $\hat{f}_X \lesssim 10^{15}$ GeV. We therefore conclude that the decays of a_λ to KK gravitons will not have a significant impact on the total widths of the a_λ.

Taken together, these results strongly suggest that intra-ensemble decays do not play a significant role in the phenomenology of bulk-axion dynamical dark matter, and can therefore be neglected. Therefore, from this point forward we will ignore intraensemble decays and identify Γ_λ, as given in Eq. (3.8), with the total width of any given a_λ.

C. Axion lifetimes across the ensemble

In Fig. 1, we show how the lifetime $\tau_\lambda \equiv 1/\Gamma_\lambda$ of an axion mass eigenstate a_λ behaves as a function of λ. The left panel shows the results for an axion with $\Lambda_G = 1$ GeV, while the right panel shows the results for an axion with $\Lambda_G = 1$ TeV. In each case, we have taken $g_G = \xi = 1$ and set $M_c = 10^{-11}$ GeV. In each of the two panels, the three solid curves correspond to three different choices of $\hat{f}_X$ for a photonic axion with $c_\gamma = 1$. The solid red curve corresponds to $\hat{f}_X = 10^8$ GeV, the solid green curve corresponds to $\hat{f}_X = 10^{12}$ GeV, and the solid blue curve corresponds to $\hat{f}_X = 10^{16}$ GeV. The dashed curves correspond to the same choices of $\hat{f}_X$ for an axion with $c_\gamma = c_g = 1$ and $c_i = 1$ for $i = \{e, \nu_e, \nu_\mu, \nu_\tau\}$. The series of kinks which are evident in each dashed curve correspond to the thresholds at $m_{\nu_2}, m_{\nu_3}, m_e,$ and m_μ above which new decay channels for a_λ open up. The sharp drop in τ_λ at around $\lambda \sim 400$ MeV is the result of the $a \to \pi^+ \pi^- \pi^0$ decay channel opening up.

One significant property of the decay rates of the a_λ in bulk-axion scenarios can be readily appreciated upon comparing the curves appearing in the two panels of Fig. 1: the total width Γ_λ is independent of m_X (and therefore independent of Λ_G) in the limit in which $\lambda \gg \pi m_X^2 / M_c$. This can also be seen from Eq. (3.8). It therefore follows that the

corresponding curves appearing in the two panels of this figure should coincide for values of λ above the threshold at which this condition is met for both of the selected values of Λ_G. Indeed, we see that this is in fact the case. The $\hat{f}_X = 10^{16}$ GeV curves coincide for nearly the entirety of the range of λ shown, since $\pi m_X^2 / M_c$ is approximately 9.95×10^{-12} GeV for $\Lambda_G = 1$ TeV and is far smaller for $\Lambda_G = 1$ GeV. The $\hat{f}_X = 10^{12}$ GeV curves, on the other hand, begin to coincide only for $\lambda \sim 10^{-3}$ GeV, which is just above the threshold $\pi m_X^2 / M_c \approx 9.95 \times 10^{-4}$ GeV for $\Lambda_G = 1$ TeV.

There is, however, an even more important message to be gleaned from Fig. 1. Note that in each panel, we have included for reference a set of horizontal, dashed lines indicating the time scales associated with the beginning of BBN ($t_{\text{BBN}} \sim 1$ s), the present age of the universe ($t_{\text{now}} \sim 4.3 \times 10^{17}$ s), and the usual limit on the lifetime τ_χ of a single decaying dark-matter candidate χ given in Eq. (1.2). *These benchmark times are absolutely critical for the survival of our dynamical dark-matter model.* Any a_λ with a lifetime that falls between t_{BBN} and τ_χ has the potential to disrupt BBN predictions for the abundances of light elements, distort the CMB to an unacceptable degree [28,36], produce too large a flux of X-ray or gamma-ray photons, etc. For this reason, the success of our dynamical dark-matter model rests upon the assumption that such a_λ
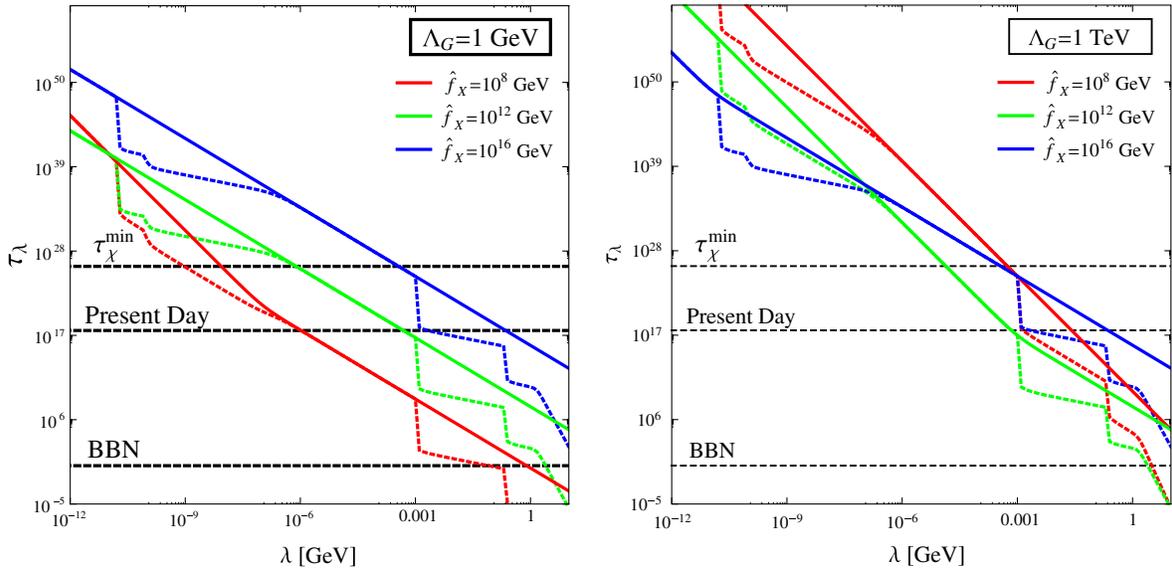

FIG. 1 (color online). The lifetime τ_λ of the axion mass eigenstate a_λ, shown as a function of its mass λ for $\Lambda_G = 1$ GeV (left panel) and an axion with $\Lambda_G = 1$ TeV (right panel). In both panels, we have set $\xi = g_G = 1$ and have chosen $M_c = 10^{-11}$ GeV—a value just above the lower bound imposed by modified-gravity experiments. The three solid curves in each panel shown correspond to different choices of $\hat{f}_X$ for a photonic axion with the coupling-coefficient assignments $c_\gamma = 1$, $c_g = 0$, and $c_i = 0$ for all fermions $i = \{e, \nu_e, \nu_\mu, \nu_\tau\}$. The solid red curve corresponds to $\hat{f}_X = 10^8$ GeV, the solid green curve corresponds to $\hat{f}_X = 10^{12}$ GeV, and the solid blue curve corresponds to $\hat{f}_X = 10^{16}$ GeV. The dashed curves correspond to the same choices of $\hat{f}_X$ for an axion with the coupling assignments $c_\gamma = c_g = 1$ and $c_i = 1$ for all i. The kinks in the curves reflect the opening up of new decay channels as λ is increased past a series of kinematic thresholds associated with decays to neutrino pairs, electron pairs, muon pairs, and $\pi^+ \pi^- \pi^0$. The horizontal lines indicate the time scales associated with the onset of BBN ($t_{\text{BBN}} \sim 1$ s), the present age of the universe ($t_{\text{now}} \sim 4.3 \times 10^{17}$ s), and the usual lower limit given in Eq. (1.2) on the lifetime τ_χ of a single-particle dark-matter candidate.

have sufficiently small relic abundances Ω_λ that these decays are harmless. *It is in this manner that lifetimes must be balanced against abundances across our dark-axion towers.*

The results in Fig. 1 also highlight another important aspect of our dynamical dark-matter ensemble, which is that at a given time t, only a fraction of the a_λ—those which have not already decayed—can contribute significantly to $\Omega_{\rm CDM}$. Therefore, since Γ_λ increases monotonically as a function of λ, there exists a maximum value $\lambda_{\rm dec}$ for which a_λ may be considered stable for particular time scale t (in the sense that $\Gamma_\lambda t < 1$) and which is potentially capable of contributing significantly to $\Omega_{\rm CDM}$. For a photonic axion, for example, the approximate form of $\lambda_{\rm dec}$ can readily be obtained in both the $y \ll 1$ and $y \gg 1$ regimes by inverting Eq. (3.1):

$$\lambda_{\rm dec} \approx \begin{cases} \left(\dfrac{g_G^4 \xi^4}{2(32\pi)^2 G_\gamma t_{\rm now}}\right)^{1/5}\left(\dfrac{\Lambda_G^{8/5}}{\hat{f}_X^{2/5} M_c^{2/5}}\right) & y \ll 1 \\[2ex] \left(\dfrac{1}{2G_\gamma t_{\rm now}}\right)^{1/3}\hat{f}_X^{2/3} & y \gg 1. \end{cases} \quad (3.14)$$

In Fig. 2, we illustrate the behavior of $\lambda_{\rm dec}$ (solid curves) for a photonic axion a function of $\hat{f}_X$ for several choices of M_c. In each case, we have fixed $\Lambda_G = 10$ GeV and set $\xi = g_G = c_\gamma = 1$. For reference, we have also included the corresponding curves for two other critical values of λ in any given axion KK tower for each M_c. These are the mass λ_0 (dashed curves) of the lightest axion mass eigenstate a_{λ_0}

and the mass $\lambda_{\rm trans}$ (dotted curves) defined in Eq. (2.28) which delineates the transition point between the small-λ and large-λ regimes. We have also included a pair of black curves indicating the asymptotic behavior of λ_0 and $\lambda_{\rm dec}$ for large M_c. Note that in this limit, mixing is negligible, and $\lambda_0 \approx m_X$ for all values of $\hat{f}_X$.

The behavior of λ_0, $\lambda_{\rm dec}$, and $\lambda_{\rm trans}$ depends primarily on the value of the mixing parameter y. When $\hat{f}_X$ is sufficiently large that $y \gg 1$ for a given value of M_c (i.e., in the lower right portion of Fig. 2), we find that $\lambda_0 \approx m_X$, as expected. At the same time, since mixing is negligible in this regime, increasing $\hat{f}_X$ results in a uniform suppression of the decay widths of all a_λ, and hence $\lambda_{\rm dec}$ increases with increasing $\hat{f}_X$ (i.e., in the upper right portion of the figure). However, as $\hat{f}_X$ decreases past the point at which the λ_0 and $\lambda_{\rm trans}$ curves intersect, we pass from the $y \gg 1$ to the $y \ll 1$ regime. In this latter regime, mixing is significant and $\lambda_0 \approx M_c/2$. At first, $\lambda_{\rm dec}$ still decreases with decreasing $\hat{f}_X$ in this regime, as those modes with $\lambda \lesssim \pi m_X^2/M_c$ continue to destabilize down the tower. However, as $\hat{f}_X$ decreases still further, to the point at which the $\lambda_{\rm dec}$ and $\lambda_{\rm trans}$ curves intersect, the a_λ in the $\lambda \lesssim \pi m_X^2/M_c$ regime begin to destabilize as well. The dependence of Γ_λ on $\hat{f}_X$ is qualitatively different for such modes, as indicated in Eq. (3.14), and consequently $\lambda_{\rm dec}$ actually begins to *increase* with decreasing $\hat{f}_X$. These observations will turn out to be critical in interpreting the results to be derived in Sec. V.
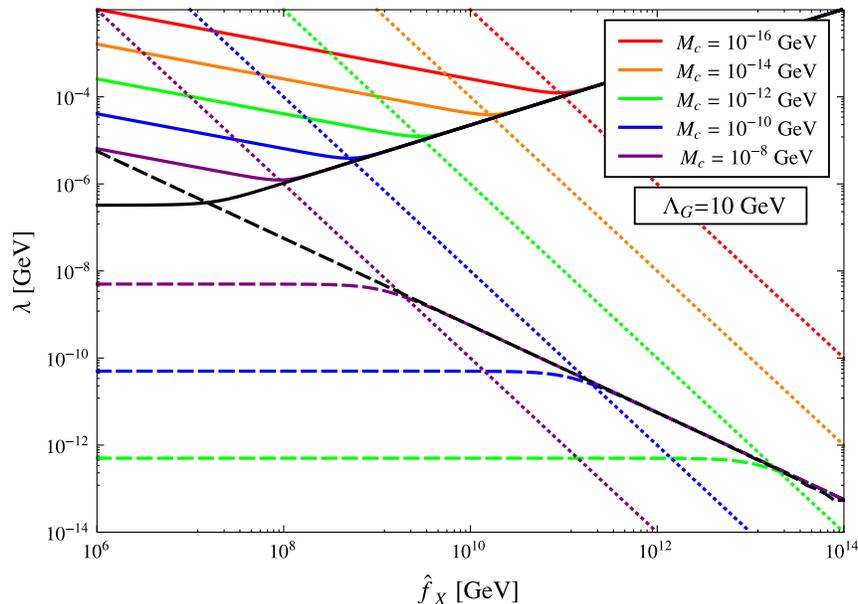

FIG. 2 (color online).   Curves showing a variety of critical values of λ as functions of $\hat{f}_X$ for $\Lambda_G = 10$ GeV and several different choices of M_c ranging from 10^{-16} GeV to 10^{-8} GeV. Each dashed curve indicates the mass λ_0 of the of the lightest axion mass eigenstate a_{λ_0} for a given choice of M_c. Each solid curve indicates the mass $\lambda_{\rm dec}$ of the heaviest a_λ for that choice of M_c which has not decayed by present time, assuming a photonic axion with $c_\gamma = 1$. Each dotted curve marks the transition point $\lambda_{\rm trans}$ between the small-λ and large-λ regimes for the same choice of M_c, as defined in in Eq. (2.28). The black curves indicate the asymptotic behavior of λ_0 and $\lambda_{\rm dec}$ for large M_c.

IV. CHARACTERIZING THE CONSTITUENTS: RELIC ABUNDANCES

In the previous section, we focused on one of the two crucial properties of the particles that constitute a dynamical dark-matter ensemble: their stability. In particular, we examined the decay widths of the various a_λ and investigated how these widths scale with λ. In this section, we focus on the other property: their abundances. We begin by establishing a consistent cosmological context in which to situate our bulk-axion theory. We then proceed to evaluate the various mechanisms through which a population of a_λ can be produced in the early universe, and explain why misalignment production is favored from the perspective of dynamical dark matter. This thereby justifies the emphasis placed on this mechanism in Ref. [13]. We then derive explicit formulas for Ω_λ, and demonstrate that the proper balance between Γ_λ and Ω_λ is indeed realized in the context of misalignment production. In the following section, we will then use these results for Γ_λ and Ω_λ to characterize the aggregate properties of the entire ensemble, such as its equation of state, its total relic abundance, and the way in which that abundance is partitioned among its constituents.

A. Standard and low-temperature reheating (LTR) cosmologies

Before embarking on a discussion of axion production in the early universe, however, we must first specify the cosmology in which that production occurs. This is particularly relevant in the context of theories involving large extra dimensions, since the properties of the early universe in such theories can differ dramatically from those which characterize the standard cosmology. For example, the presence of substantial energy density in the bulk can alter the expansion rate of the universe in a significant way [37], and late decays of KK excitations of the graviton (or other bulk fields) can disrupt BBN, produce visible distortions in the diffuse photon spectrum, etc. For this reason, such scenarios must obey stringent constraints [28] on the so-called "normalcy temperature" T_*: the temperature below which the universe is effectively four-dimensional, in the sense that the bulk is essentially empty of energy density and the radii of the extra dimensions can be regarded as fixed. These bounds come from a diverse array of considerations and leave a very narrow window of 4 MeV $\lesssim$ $T_* \lesssim$ 30 MeV for this normalcy temperature. The most attractive solution for arranging such a value of T_* is to posit a very late period of cosmic inflation, precipitated by a brane-localized inflaton [28], with a reheating temperature $T_{\mathrm{RH}} \simeq T_*$.

If T_{RH} is indeed to be identified with T_*, then the universe must be described not by the standard cosmology, but by an alternative framework commonly dubbed the low-temperature reheating (LTR) cosmology [38]. In the LTR framework, inflation occurs very late, and the energy density of the universe remains dominated by coherent oscillations of the inflaton field ϕ. Such oscillations behave like massive matter down to very low temperatures—potentially as low as a few MeV. Consequently, the universe undergoes an additional epoch of matter domination (MD) at early times, which ends only once decays of ϕ into SM fields in the radiation bath have depleted the energy density ρ_ϕ to the point that $\rho_\phi = \rho_{\mathrm{rad}}$. This condition defines the reheating temperature T_{RH}, which is determined solely by the decay width Γ_ϕ of the inflaton:

$$T_{\mathrm{RH}} \equiv \left[\frac{90}{\pi^2 g_*(T_{\mathrm{RH}})} \right]^{1/4} \sqrt{\Gamma_\phi M_P}, \qquad (4.1)$$

where $g_*(T)$ denotes the effective number of massless, interacting degrees of freedom at temperature T. At temperatures $T \lesssim T_{\mathrm{RH}}$ the universe is effectively radiation dominated (RD), and maps onto the standard cosmology. This modification implies that the relationship between time and temperature in the LTR cosmology is quite different from that obtained in the standard cosmology at early times. In particular, one finds that in any universe which underwent a period of cosmic inflation immediately prior to the RD era, the relationship between time and temperature is given by

$$t = \begin{cases} \sqrt{\frac{45}{2\pi^2}} \dfrac{g_*^{1/2}(T_{\mathrm{RH}})}{g_*(T)} \dfrac{T_{\mathrm{RH}}^2 M_P}{T^4} & T_{\mathrm{max}} \gtrsim T \gtrsim T_{\mathrm{RH}} \\[2ex] \sqrt{\frac{45}{2\pi^2}} g_*^{-1/2}(T) \dfrac{M_P}{T^2} & T_{\mathrm{RH}} \gtrsim T \gtrsim T_{\mathrm{MRE}} \\[2ex] \sqrt{\frac{45}{2\pi^2}} g_*^{-1/8}(T_{\mathrm{MRE}}) g_*^{-3/8}(T) \dfrac{M_P}{T_{\mathrm{MRE}}^{1/2} T^{3/2}} & T \lesssim T_{\mathrm{MRE}}, \end{cases}$$

$$(4.2)$$

where $T_{\mathrm{MRE}} \sim \mathcal{O}(\mathrm{eV})$ is the temperature associated with the usual matter/radiation transition, at which the energy density of the universe once again becomes dominated by the contributions from dark and baryonic matter. (For our purposes, it will be sufficient to approximate the present, ΛCDM universe as matter-dominated.) In the standard cosmology, T_{RH} is high enough that the universe will be radiation-dominated at all time scales relevant to axion dynamics, all the way down to the time scale t_{MRE} associated with this transition. By contrast, in a LTR cosmology with $T_{\mathrm{RH}} \sim \mathcal{O}(\mathrm{MeV})$, much of the relevant dynamics will occur while the universe is still dominated by coherent oscillations of the inflaton field. Such a cosmological modification can have profound effects on axion dynamics, even in the case of a four-dimensional axion [39,40].

Since our model necessarily involves large extra dimensions, constraints on T_* would seem to require that the universe be described by the LTR cosmology, rather than the standard cosmology, at early times. Indeed, for this reason, we shall adopt such an LTR cosmology in what follows. This will actually turn out to be advantageous, as the relic abundance of a light scalar generated via misalignment production can differ substantially between the two

cosmologies [40]. To better facilitate comparison between the results obtained in these two cosmologies, we therefore find it instructive to present results for each in parallel. The reader should keep in mind, however, that the results obtained for an LTR cosmology with $4 \text{ MeV} \lesssim T_* \lesssim 30 \text{ MeV}$ should be taken, in some sense, as the "true" ones, given the constraints on T_*.

In this paper, we shall be interested in values of the confinement scale Λ_G which range between roughly 10 MeV and 100 TeV. When operating in the standard cosmology, we shall assume that $\Lambda_G \ll T_{\text{RH}}$, so that confinement takes place *within the RD era.* Conversely, when operating within the LTR cosmology, we will assume that $\Lambda_G > T_{\text{RH}}$, so that confinement takes place *before or during reheating*, when the universe is dominated by coherent oscillations of the inflaton field. Note, however, that for smaller confinement scales $\Lambda_G < T_{\text{RH}}$, the results for the LTR cosmology are identical to those for the standard cosmology.

B. Axion production mechanisms

Having now set the cosmological tableaux, let us begin our discussion of axion production in the early universe. A number of production mechanisms can, in principle, contribute significantly to the axion relic density. In generic models, three such mechanisms typically tend to provide the dominant contribution to the relic density of any individual axion field. One of these mechanisms is thermal production via the interactions of the a_λ with the SM fields in the radiation bath. The other two are nonthermal in nature and can generate a population of cold axions. These are misalignment production and production from the decays of topological defects (in particular, cosmic strings) associated with the breaking of the global $U(1)_X$ symmetry.

We begin with a discussion of thermal production. A number of processes can contribute appreciably to thermal axion production in the early universe, depending of course on the magnitudes of the couplings between the axion in question couples to the SM particles. Among hadronic processes, $qg \to qa_\lambda$, $q\bar{q} \to ga_\lambda$, $gg \to ga_\lambda$, etc., dominate for $T \gtrsim \Lambda_{\text{QCD}}$, while pion-axion conversion off nuclei (including all processes of the form $N\pi \to N'a_\lambda$, where $N, N' = \{n, p\}$) and the purely pionic process $\pi\pi \to \pi a_\lambda$ dominate at lower temperatures. Since we are interested in values of M_c which are far below T_{RH}, it follows that $T \gg \lambda$ for a large number of the a_λ for at least some of the post-inflationary era. These a_λ can therefore be considered effectively massless at such temperatures. In this massless limit, the rate for each of the axion-production processes enumerated above (except for the inverse-decay process, which is generally subleading) takes the rough, parametric form

$$\Gamma \propto \frac{T^3}{\hat{f}_X^2}(\tilde{\lambda}^2 A_\lambda)^2. \tag{4.3}$$

In other words, since kinematic distinctions between states for which $T \gg \lambda$ are unimportant, the dependence of their production rates on λ occurs primarily through the coupling factor $\tilde{\lambda}^4 A_\lambda^2$. This factor too is effectively independent of λ when $\lambda \gtrsim \pi m_X^2/M_c$; hence, at a given temperature T, any a_λ with a mass λ in the range $\pi m_X^2/M_c \lesssim \lambda \lesssim T$ will be produced at essentially the same rate. Moreover, production rates actually *increase* with increasing λ for those modes with masses $\lambda \lesssim \pi m_X^2/M_c$. This means that at a given temperature T, the heavier a_λ in this mass range will actually be produced from the thermal bath at an equal or higher rate than the lighter a_λ—at least until λ becomes comparable with T, and their production rates are Boltzmann suppressed. Since Γ_λ also increases with λ, the less stable states will be thermally produced at an equal or higher rate than the more stable ones. Such a relationship is clearly undesirable in models of dynamical dark matter. Moreover, the majority of light axions produced through interactions with the thermal bath would be relativistic at the time of production, and therefore not cold.

From these considerations, we conclude that if an ensemble of axion KK modes is to constitute the majority of the dark-matter relic density, thermal production must contribute only a negligible fraction of the total relic abundance of each a_λ, with the remainder of that abundance generated through nonthermal means. This requirement places a nontrivial constraint on scenarios of this sort, a detailed analysis of which appears in Ref. [32]. We will defer the discussion of how this constraint restricts the parameter space of our model until Sec. VI. For the moment, we simply note that this constraint exists, and proceed to discuss nonthermal mechanisms for axion production. Note, however, that in traditional models of KK dark-matter (either single-component [6] or multicomponent [5]), in which the dark-matter candidates are stable, thermal freeze-out can be a viable production method for generating relic abundances.

One method in which a nonthermal population of axions may be generated in the early universe is production via the decay of cosmic strings associated with the broken global $U(1)_X$ symmetry. However, this mechanism can contribute significantly to axion production only if $H_I \gtrsim f_X$, where H_I is the value of the Hubble parameter during inflation, so that those cosmic strings are not inflated away. Since the value of H_I is relatively unconstrained, and since astrophysical and cosmological constraints will turn out to require $\hat{f}_X$ to be quite large, in the rest of this paper, we shall assume that $H_I \ll f_X$. We will therefore not consider axion production from $U(1)_X$ string decay. However, it should be noted that in other scenarios (or in other regions of parameter space), axions produced via cosmic-string decay could have important phenomenological consequences in dynamical dark-matter models, and this production mechanism therefore deserves further study.

Finally, we turn to nonthermal axion production via the misalignment mechanism. As we shall see, this turns out to be the most promising axion-production mechanism from the perspective of dynamical dark matter. The basis of this mechanism is that at temperatures $T \gg \Lambda_G$, the instanton-induced contribution to the axion potential in Eq. (2.18) effectively vanishes. This implies that the only contribution to the axion mass matrix at such high temperatures are the diagonal contributions from the KK masses: no mixing occurs, and consequently the mass eigenstates are merely the KK eigenstates a_n. While the potential for each a_n with $n \neq 0$ is therefore nonvanishing, due to the presence of the KK masses, and is minimized at $a_n = 0$, the potential for the zero mode a_0 vanishes. In the absence of a potential for a_0, there is no preferred vacuum expectation value $\langle a_0 \rangle$ which minimizes $V(a_0)$; indeed, any $\langle a_0 \rangle \lesssim \hat{f}_X$ is as good as any other. This means that when the $U(1)_X$ symmetry is broken, the value of $\langle a_0 \rangle$ within a given domain is essentially arbitrary. Thus, immediately after this phase transition occurs, one would expect to find a set of domains, each with a different homogeneous background value for the axion field, which would generically be expected to be $\mathcal{O}(\hat{f}_X)$, but could in principle be smaller. Our ignorance of this initial value of $\langle a_0 \rangle$ is commonly parameterized by a "misalignment angle" θ, so that the initial conditions at the time at which the $U(1)_X$ symmetry is broken can be written as [26]

$$\langle a_0 \rangle = \theta \hat{f}_X, \qquad \langle a_n \rangle = 0 \quad \text{for } n \neq 0. \tag{4.4}$$

Indeed, as was noted in Ref. [13], this initial condition follows from $U(1)_X$ invariance, which manifests itself here in the form of a five-dimensional shift symmetry under which $a \to a + c$, where c is a constant.

At lower temperatures, however, the situation changes, as instanton effects generate a brane mass $m_X(T)$ for the bulk axion. Here, we write $m_X(T)$ rather than m_X in order to emphasize that this instanton-induced mass term is temperature dependent, and reserve the symbol m_X (without the argument) to refer to the constant, late-time (i.e., low-temperature) value of $m_X(T)$. Assuming that the instantons associated with the group G behave analogously to QCD instantons, Eq. (2.5) implies that $m_X(T)$ scales roughly like $(\Lambda_G/T)^4$. Thus, when $T \sim \Lambda_G$, the off-diagonal terms in the mass matrix $\mathcal{M}^2_{mn}(T)$ become appreciable, and the a_n are no longer mass eigenstates. In this regime, the equations for a_n form a coupled system [26], with the evolution of each such field governed by an equation of the form

$$\ddot{a}_n + 3H\dot{a}_n + \sum_\lambda \Gamma_\lambda(T)U_{\lambda n}(T)\dot{a}_n + \sum_{m=0}^{\infty} \mathcal{M}^2_{nm}(T)a_m = 0, \tag{4.5}$$

where $U_{\lambda n}(T)$ denotes the unitary matrix in Eq. (2.23), a dot denotes a derivative with respect to the time t, and H is the Hubble parameter. Note that since the mass eigenvalues $\lambda(T)$, decay widths $\Gamma_\lambda(T)$, and even the rotation matrix $U_{\lambda n}(T)$ itself all depend on m_X, these quantities all implicitly depend on temperature, and hence on t.

During any period in which $\Gamma_\lambda(T)$ and $\lambda(T)$ vary appreciably in time, it is in general not possible to write down an exact, closed-form solution to the coupled system in Eq. (4.5). Fortunately, however, $m_X(T)$ can be regarded as effectively constant during most of the history of the universe. The only exception occurs at temperatures around $T \sim \Lambda_G$, where $m_X(T)$ rapidly rises from a negligible initial value to the asymptotic value it attains at $T \ll \Lambda_G$. At all other times, $m_X(T)$ is well approximated either by zero or by m_X, and the system of equations therefore decouples in the mass-eigenstate basis. In this basis, the time-evolution of each field a_λ is governed by an equation of the form

$$\ddot{a}_\lambda + \frac{\kappa}{t}\dot{a}_\lambda + \Gamma_\lambda \dot{a}_\lambda + \lambda^2 a_\lambda = 0. \tag{4.6}$$

In arriving at this expression, we have used the fact that within an RD or MD era, H is approximately given by the relation $3H \approx \kappa/t$, where

$$\kappa \equiv \begin{cases} 3/2 & \text{in RD} \\ 2 & \text{in MD.} \end{cases} \tag{4.7}$$

Exact, closed-form solutions for a_λ and $\dot{a}_\lambda$, given an evolution equation of this form, *do* exist, and we present these solutions in the Appendix. It may be observed, however, that Eq. (4.6) is simply the equation of motion for a damped harmonic oscillator with a time-dependent damping term. As we shall see at the end of this section, it turns out that $\Gamma_\lambda \ll 3H$ at the time when $\lambda \approx 3H/2$ for all a_λ. It then follows that the solutions for each a_λ can be divided into two regimes, depending on the relationship between λ and H at any given time t. When $\lambda \lesssim 3H/2$, a_λ does not oscillate, and therefore its energy density scales approximately like vacuum energy. By contrast, when $\lambda \gtrsim 3H/2$, a_λ oscillates coherently around the minimum of its potential, with oscillations damped by a "friction" term with coefficient $(3H + \Gamma_\lambda)$.

Despite the fact that $m_X(T)$ is effectively constant both well before and well after the time scale t_G at which $T = \Lambda_G$, the nontrivial dynamics of the a_λ at $t \sim t_G$ certainly can play a crucial role in establishing the initial relic abundances for these fields. Nevertheless, while such a time-dependence can indeed have a significant quantitative impact on the results of relic-density calculations in certain cases, as has been shown to be the case with a standard, four-dimensional axion [22], a great deal of information can be obtained by working in the "rapid-turn-on" approximation, in which we approximate

$$m_X(t) = m_X\Theta(t - t_G). \tag{4.8}$$

In this approximation, the procedure for calculating the background value of each a_λ at any time t is clear [13].

At early times, when $t < t_G$, $m_X(T) = 0$, and hence the a_n remain fixed at the initial values given in Eq. (4.4). At $t = t_G$, $m_X(T)$ immediately assumes its constant, nonzero, late-time value m_X, and each of the mass eigenstates a_λ acquires a background value proportional to its overlap with a_0:

$$\langle a_\lambda \rangle = \theta \hat{f}_X A_\lambda, \qquad \langle \dot{a}_\lambda \rangle = 0 \quad \text{at } t = t_G. \qquad (4.9)$$

Even though all of the a_λ acquire background values at $t = t_G$, only those fields for which $\lambda \gtrsim 3H(t_G)/2$ begin oscillating immediately at the time of this phase transition. As discussed above, all other, lighter a_λ will begin oscillating later, once the $\lambda \gtrsim 3H/2$ threshold is crossed. The time t_λ at which a given a_λ begins to oscillate is therefore given by

$$t_\lambda \equiv \max \left\{ \frac{\kappa_\lambda}{2\lambda}, t_G \right\}, \qquad (4.10)$$

where κ_λ is the value of κ corresponding to the epoch during which this oscillation begins. At times $t_G < t \lesssim t_\lambda$, a given a_λ continues to behave like vacuum energy rather than like matter, and thus properly contributes not to Ω_{CDM}, but to the dark-energy abundance. (Note that this definition of t_λ is slightly different from the one given in Ref. [13], where t_λ was simply defined as κ_λ/λ, regardless of its relationship to t_G.)

For any given a_λ, however, the relevant quantity for dark-matter phenomenology is not the value of a_λ itself, but its energy density ρ_λ, which is related to a_λ and $\dot{a}_\lambda$ by the relation

$$\rho_\lambda = \frac{1}{2} [\dot{a}_\lambda^2 + \lambda^2 a_\lambda^2]. \qquad (4.11)$$

At early times, when $t < t_G$, only a_0 has a nonzero background value in the rapid-turn-on approximation, and since this field is massless, its energy density vanishes. At $t = t_G$, however, Eq. (4.9) implies that each field acquires an initial energy density

$$\rho_\lambda(t_G) = \frac{1}{2} \theta^2 \hat{f}_X^2 \lambda^2 A_\lambda^2. \qquad (4.12)$$

Since a_λ remains effectively constant until $t = t_\lambda$ for any field for which $t_\lambda > t_G$, we also see that $\rho_\lambda(t_\lambda) = \rho_\lambda(t_G)$. This implies that the energy density stored in such a field behaves like vacuum energy until $t = t_\lambda$, at which point the field begins to oscillate coherently around the minimum of its potential. The energy density stored in such oscillations, as is well known, scales like massive matter. At late times $t \gg t_\lambda$, when the time scale associated with these oscillations becomes rapid compared to the time scale over which the amplitude of a_λ changes appreciably, the virial approximation implies that $\rho_\lambda \approx \langle \dot{a}_\lambda^2 \rangle$, where $\langle \dot{a}_\lambda^2 \rangle$ denotes the average of $\dot{a}_\lambda^2$ over one cycle of oscillation. In this regime, one finds that

$$\rho_\lambda(t) = \rho_\lambda(t_G) \left(\frac{t_\lambda}{t} \right)^{\kappa_\lambda} e^{-\Gamma_\lambda(t - t_G)} \qquad (4.13)$$

for each ρ_λ during the epoch in which oscillation began, with $\rho_\lambda(t_G)$ given in Eq. (4.12). Computing ρ_λ during subsequent epochs is simply a matter of applying Eq. (4.13) iteratively with the appropriate boundary conditions for ρ_λ at the transition points at which κ changes.

All that remains, then, in order to specify the energy density ρ_λ associated with a given a_λ for any particular choice of model parameters is to determine the time scales t_G and t_λ as a function of those parameters. Indeed, the results for ρ_λ clearly depend sensitively both on the epoch during which abundances are established, and the epoch during which oscillation begins. In principle, t_G could fall within the RD era, during the usual MD era, or during reheating; likewise, t_λ could occur at any time at or after t_G. However, as will be shown in Sec. V, $\Lambda_G \gg 10$ MeV is required in order for our ensemble of a_λ to yield a relic abundance on the order of Ω_{CDM}. For such values of Λ_G, $t_G \gg t_{\mathrm{MRE}}$. Moreover, for $\mathcal{O}(1)$ values of g_G and ξ, $t_\lambda \ll t_{\mathrm{MRE}}$ as well. Therefore, when discussing the standard cosmology, we will focus on the case in which $t_G \leq t_\lambda < t_{\mathrm{MRE}}$ for all a_λ. Furthermore, when operating within the context of the LTR cosmology, we will implicitly assume that $t_G \leq t_\lambda < t_{\mathrm{RH}}$ for all a_λ, so that all fields begin oscillating while the energy density of the universe is still dominated by coherent oscillations of the inflaton field. This is justified by the fact that $t_{\mathrm{RH}} \sim 10^{-1} - 10^{-4}$ s for reheating temperatures T_{RH} within the phenomenologically allowed window $4 \text{ MeV} \lesssim T_* \lesssim 30 \text{ MeV}$. In conjunction with the experimental bound on M_c given in Eq. (3.4), this implies that $t_\lambda \lesssim t_{\mathrm{RH}}$ for all a_λ in any given tower, unless $\Lambda_G \lesssim 10$ MeV. In summary, we shall therefore assume that

$$
\begin{aligned}
\text{standard cosmology:} \quad & t_{\mathrm{RH}} < t_G \leq t_\lambda < t_{\mathrm{MRE}} \\
\text{LTR cosmology:} \quad & t_P < t_G \leq t_\lambda < t_{\mathrm{RH}}
\end{aligned} \qquad (4.14)
$$

in what follows, where t_P is the Planck time. Note that in the standard cosmology, t_{RH} is assumed to be so early that all t_G of interest will easily satisfy the lower bound. Also note that in the LTR cosmology, modes for which t_λ occurs during or prior to inflation will inflate away and therefore carry zero abundance at present time.

C. Axion relic abundances

We now provide explicit expressions for ρ_λ in both the standard and LTR cosmologies. We begin by considering the case of the standard cosmology, in which $\kappa = 3/2$ at all relevant time scales prior to matter-radiation equality, and $\kappa = 2$ after the transition to matter-domination at t_{MRE}. In this cosmological framework, given the regimes for t_G and t_λ specified in Eq. (4.14), it therefore follows that

$$\rho_\lambda^{\mathrm{Std}}(t) = \frac{1}{2}\theta^2 \hat{f}_X^2 \lambda^2 A_\lambda^2 e^{-\Gamma_\lambda(t-t_G)} \times \begin{cases} \left(\frac{t_\lambda}{t}\right)^{3/2} & t_\lambda \lesssim t \lesssim t_{\mathrm{MRE}} \\[2mm] \left(\frac{t_\lambda^{3/2} t_{\mathrm{MRE}}^{1/2}}{t^2}\right) & t \gtrsim t_{\mathrm{MRE}}. \end{cases}$$

(4.15)

By contrast, in the LTR cosmology, the energy of the universe remains dominated by coherent oscillations of the inflaton field from the end of inflation until a very late time $t_{\mathrm{RH}} \sim 1$ s. Thus, if the axion fields begin oscillating at a time $t_\lambda < t_{\mathrm{RH}}$, as specified in Eq. (4.14), we initially have $\kappa = 2$, followed by a transition at t_{RH} to the usual RD era, in which $\kappa = 3/2$. This signifies that in the LTR cosmology, we have

$$\rho_\lambda^{\mathrm{LTR}}(t) = \frac{1}{2}\theta^2 \hat{f}_X^2 \lambda^2 A_\lambda^2 e^{-\Gamma_\lambda(t-t_G)}$$

$$\times \begin{cases} \left(\frac{t_\lambda}{t}\right)^2 & t_\lambda \lesssim t \lesssim t_{\mathrm{RH}} \\[2mm] \left(\frac{t_\lambda^2}{t_{\mathrm{RH}}^{1/2} t^{3/2}}\right) & t_{\mathrm{RH}} \lesssim t \lesssim t_{\mathrm{MRE}} \\[2mm] \left(\frac{t_\lambda^2 t_{\mathrm{MRE}}^{1/2}}{t^2 t_{\mathrm{RH}}^{1/2}}\right) & t \gtrsim t_{\mathrm{MRE}}. \end{cases}$$

(4.16)

In other words, Eq. (4.16) replaces Eq. (4.15) in the context of the LTR cosmology, in which the usual RD era is preceded by an initial period of matter domination. It is worth emphasizing here that the value of t_λ for a given mass eigenvalue λ will, in general, differ between the two cosmologies, due to the differing relationship between H and t at times $t \lesssim t_{\mathrm{RH}}$.

Comparing Eqs. (4.15) and (4.16), we see that the cosmological context in which the axion fields evolve can have a potentially dramatic effect on the late-time results for the various ρ_λ. However, for any given mode, Eq. (4.10) implies that the magnitude $\mathcal{E}_{\mathrm{LTR}}$ of that suppression depends on the relationship between λ and t_G. Comparing Eqs. (4.15) and (4.16), we find that the axion energy densities are suppressed in the LTR cosmology, relative to the standard cosmology, by a factor

$$\mathcal{E}_{\mathrm{LTR}} \equiv \frac{\rho_{\mathrm{LTR}}}{\rho_{\mathrm{Std}}}$$

$$\approx \begin{cases} \frac{g_*^{5/4}(T_{\mathrm{RH}})}{g_*^{5/4}(\Lambda_G)} \left(\frac{T_{\mathrm{RH}}}{\Lambda_G}\right)^5 & \lambda \geq 1/t_G \\[3mm] 4\left(\frac{2}{3}\right)^{3/2} \left(\frac{2\pi^2}{45}\right)^{1/4} g_*^{1/4}(T_{\mathrm{RH}}) \frac{T_{\mathrm{RH}}}{M_P^{1/2}\lambda^{1/2}} & \lambda < 1/t_G. \end{cases}$$

(4.17)

These results imply that the energy-density contributions from those modes which begin oscillating at t_G are suppressed in the LTR cosmology, relative to their value in the standard cosmology, to a greater degree than the contributions from those modes which begin oscillating later. These results are analogous to those obtained in Ref. [40] for a standard, four-dimensional axion, where the suppression factor is referred to as $V^{\mathrm{LTR}}/V^{\mathrm{Std}}$.

It is also possible (and indeed when Λ_G is large, more or less inevitable) that in the LTR cosmology, a great many of the heavier a_λ will begin to oscillate either during or prior to the end of inflation. Since the scale factor R grows exponentially during this epoch, the energy density in stored any such mode, which scales like $\rho_\lambda \propto R^{-3}$, will effectively be diluted into irrelevance by this rapid expansion. As long as $\Lambda_G \lesssim T_{\mathrm{max}}$, i.e., as long as confinement occurs only after inflation is over, no energy density is stored in these modes during inflation. As a result, the energy density is given by Eq. (4.16) as usual. However, if confinement occurs before or during inflation, ρ_λ for any mode for which $\lambda > 3H_I/2$ will be exponentially damped by Hubble dilution, and it is therefore reasonable to take $\rho_\lambda = 0$ for any such mode.

From the results in Eqs. (4.15) and (4.16), it is straightforward to obtain the relic abundance $\Omega_\lambda \equiv \rho_\lambda/\rho_{\mathrm{crit}}$ for each a_λ. Let us begin by addressing those modes for which $t_\lambda = t_G$. Since the critical density for a flat universe is given by $\rho_{\mathrm{crit}} = 3H^2 M_P^2$, we find that in the rapid-turn-on approximation, the contribution to the dark-matter relic abundance from each such mode at a given time t in the standard cosmology is

$$\Omega_\lambda^{\mathrm{Std}} = 3\left(\frac{\theta\hat{f}_X m_X}{M_P}\right)^2 t_G^{3/2} \left[1 + \frac{\lambda^2}{m_X^2} + \frac{\pi^2 m_X^2}{M_c^2}\right]^{-1} e^{-\Gamma_\lambda(t-t_G)}$$

$$\times \begin{cases} \frac{4}{9} t^{1/2} & t \lesssim t_{\mathrm{MRE}} \\[2mm] \frac{1}{4} t_{\mathrm{MRE}}^{1/2} & t \gtrsim t_{\mathrm{MRE}}. \end{cases}$$

(4.18)

By contrast, in the LTR cosmology, the corresponding result is

$$\Omega_\lambda^{\mathrm{LTR}} = 3\left(\frac{\theta\hat{f}_X m_X}{M_P}\right)^2 t_G^2 \left[1 + \frac{\lambda^2}{m_X^2} + \frac{\pi^2 m_X^2}{M_c^2}\right]^{-1} e^{-\Gamma_\lambda(t-t_G)}$$

$$\times \begin{cases} \frac{1}{4} & 1/\lambda \lesssim t \lesssim t_{\mathrm{RH}} \\[2mm] \frac{4}{9}\left(\frac{t}{t_{\mathrm{RH}}}\right)^{1/2} & t_{\mathrm{RH}} \lesssim t \lesssim t_{\mathrm{MRE}} \\[2mm] \frac{1}{4}\left(\frac{t_{\mathrm{MRE}}}{t_{\mathrm{RH}}}\right)^{1/2} & t \gtrsim t_{\mathrm{MRE}}. \end{cases}$$

(4.19)

For the rest of the a_λ (i.e., those for which $t_\lambda \gtrsim t_G$), the corresponding results in the context of the standard cosmology are

$$\Omega_\lambda^{\text{Std}} = 3\left(\frac{3}{4}\right)^{3/2}\left(\frac{\theta\hat{f}_X m_X}{M_P}\right)^2 \lambda^{-3/2}$$

$$\times \left[1 + \frac{\lambda^2}{m_X^2} + \frac{\pi^2 m_X^2}{M_c^2}\right]^{-1} e^{-\Gamma_\lambda(t-t_G)}$$

$$\times \begin{cases} \frac{4}{9} t^{1/2} & t \lesssim t_{\text{MRE}} \\ \frac{1}{4} t_{\text{MRE}}^{1/2} & t \gtrsim t_{\text{MRE}}, \end{cases} \quad (4.20)$$

whereas in the context of the LTR cosmology, we instead have

$$\Omega_\lambda^{\text{LTR}} = 3\left(\frac{\theta\hat{f}_X m_X}{M_P}\right)^2 \lambda^{-2}\left[1 + \frac{\lambda^2}{m_X^2} + \frac{\pi^2 m_X^2}{M_c^2}\right]^{-1} e^{-\Gamma_\lambda(t-t_G)}$$

$$\times \begin{cases} \frac{1}{4} & 1/\lambda \lesssim t \lesssim t_{\text{RH}} \\ \frac{4}{9}\left(\frac{t}{t_{\text{RH}}}\right)^{1/2} & t_{\text{RH}} \lesssim t \lesssim t_{\text{MRE}} \\ \frac{1}{4}\left(\frac{t_{\text{MRE}}}{t_{\text{RH}}}\right)^{1/2} & t \gtrsim t_{\text{MRE}}. \end{cases} \quad (4.21)$$

Note that this is valid only for those a_λ for which $\lambda \lesssim 3H_I/2$. For those modes with $\lambda \gtrsim 3H_I/2$, as discussed above, effectively $\Omega_\lambda = 0$ due to Hubble dilution during inflation. During periods of matter domination, we see that, ρ_λ and ρ_{crit} scale identically with time, and consequently Ω_λ remains constant in each of these expressions. During periods of radiation domination, on the other hand, ρ_λ falls faster with time than ρ_{crit}, and Ω_λ grow like $t^{1/2}$.

Note that for $t = t_{\text{now}}$, Eqs. (4.18) through (4.21) take the forms specified in Eqs. (54) through (60) in Ref. [13], where the brane mass m appearing in these equations is identified with m_X for a bulk axion. In other words, the distinction between the Ω_λ expressions for the standard and LTR cosmologies given here is tantamount to identifying the era in which the initial abundances are established in Ref. [13].

Since we will be primarily interested in scenarios in which the dark-matter relic abundance Ω_{CDM} receives contributions from a large number of different axion mass eigenstates, it is of critical importance to determine precisely how Ω_λ scales with λ, as this describes the relative contributions to the total relic abundance Ω_{CDM} from the various states in the tower. Fortunately, in the rapid-turn-on approximation, the dependence of Ω_λ on λ stems from only three factors, as discussed in Ref. [13]. The first, of course, is the choice of cosmology. The second is how λ compares to the mass scales m_X and $\pi m_X^2/M_c$, which determines the overlap between a_λ and the KK zero mode a_0, and therefore the initial displacement of a_λ at $t = t_\lambda$. For those modes for which $\lambda \gtrsim \max\{m_X, \pi m_X^2/M_c\}$, the λ^2/m_X^2 term dominates in the factor in brackets appearing in each of Eqs. (4.15) and (4.16), and Ω_λ acquires a factor of λ^{-2}. By contrast, for those modes for which $\lambda \lesssim \max\{m_X, \pi m_X^2/M_c\}$, the constant terms dominate, and Ω_λ acquires no such dependence.

The third factor which determines how Ω_λ scales with λ in the rapid-turn-on approximation is whether or not the axion mode in question begins oscillating at t_G, or at some later time. Indeed, in Ref. [13], these cases were referred to, respectively, as the "instantaneous" and "staggered" turn-on regimes. Comparing the expressions valid for $t_\lambda = t_G$ in Eqs. (4.18) and (4.19) with those valid for $t_\lambda > t_G$ in Eqs. (4.20) and (4.21), we see that the Ω_λ for those a_λ which begin oscillating after the confining phase transition takes place acquire an additional dependence on λ. The precise form of this dependence depends on the cosmological context within which the model is embedded: in the standard cosmology, it is $\lambda^{-3/2}$; in the LTR cosmology, it is λ^{-2}. Physically, this factor stems from the fact that prior to the time it begins to oscillate coherently, the energy density in any given a_λ remains constant, whereas after oscillation begins, it scales like massive matter. Therefore, the later a given mode begins to oscillate, the longer the energy density stored in that mode will remain unaffected by cosmic expansion, and therefore the larger the present-day value of Ω_λ will be.

In order to illustrate the implications of these effects, in each of the panels of Figs. 3 and 4 we track the evolution of Ω_λ for a representative sample of a_λ within a given theory from t_G to present time. The curves shown in Fig. 3 reflect typical results which arise in the context of the standard cosmology. The left panel corresponds to a scenario with a small confinement scale $\Lambda_G = 1$ MeV, a moderate value $\hat{f}_X = 10^9$ GeV for the effective four-dimensional $U(1)_X$-breaking scale, and a small misalignment angle $\theta \approx 0.04$. The right panel corresponds to the opposite case: a scenario in which $\Lambda_G = 1$ TeV, with $\hat{f}_X \approx 2 \times 10^7$ GeV and $\theta = 1$. In each scenario, we have taken $\xi = g_G = 1$, and set the compactification scale to be $M_c = 10^{-11}$ GeV. The curves shown in each panel (from top to bottom) correspond to the lightest two values of λ, here referred to as λ_0 and λ_1, corresponding to that particular choice of $\hat{f}_X$ and M_c ($\lambda_0 \approx 5 \times 10^{-12}$ GeV and $\lambda_1 \approx 2 \times 10^{-11}$ GeV for the right panel; $\lambda_0 \approx 6 \times 10^{-17}$ GeV and $\lambda_1 \approx 10^{-11}$ GeV for the left panel), along with several additional values of λ, including $\lambda = \{10^{-10}, 10^{-8}, 10^{-6}, 10^{-4}, 10^{-2}\}$ GeV. It should be noted that in the right panel, all curves shown lie essentially on top of one another, and are thus not individually apparent. It is worth remarking that $y \gg 1$ in the scenario depicted in the left panel, while $y \ll 1$ in the scenario depicted in the right panel. Also shown are a horizontal, dot-dashed line indicating the value for Ω_{CDM} observed by the WMAP experiment [2], as quoted in Eq. (1.1), and a pair of vertical lines indicating the positions of t_{RH} and t_{MRE}.

Both of these scenarios yield a total dark-matter relic abundance which is consistent with the WMAP results given in Eq. (1.1). However, it should be noted that the parameter assignments used for these figures have been chosen exclusively for purposes of illustration. In Sec. VI,
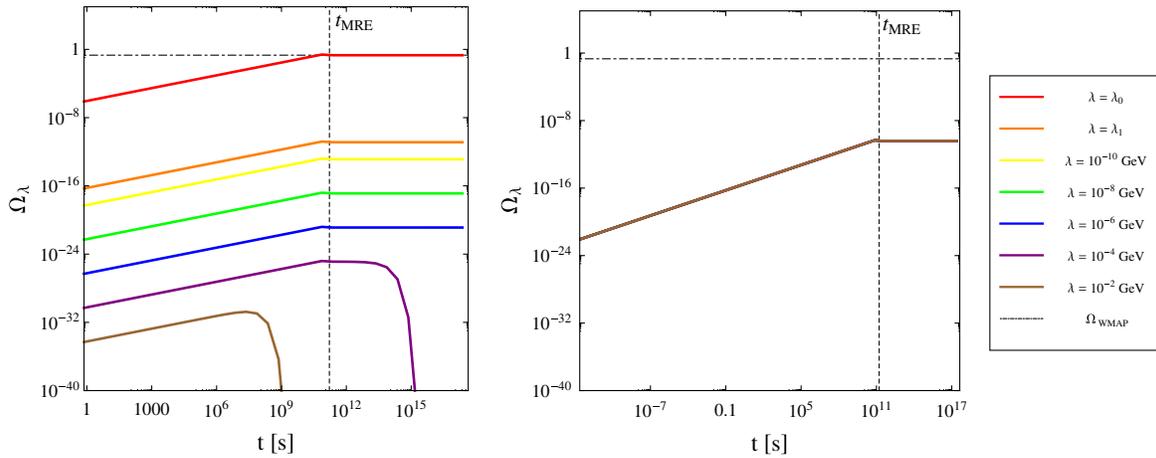
FIG. 3 (color online).   The individual relic abundances Ω_λ associated with a variety of different a_λ in the standard cosmology, shown as a function of time t, for two scenarios with different values of the parameters $\hat{f}_X$, Λ_G, and θ. The left panel corresponds to a choice of $\hat{f}_X \approx 10^9$, $\Lambda_G = 1$ MeV, and $\theta \approx 0.04$, while the right panel corresponds to a choice of $\hat{f}_X \approx 2 \times 10^7$ GeV, $\Lambda_G = 1$ TeV, and $\theta = 1$. In both cases, we have taken $\xi = g_G = 1$, with $M_c = 10^{-11}$ GeV. The range of t displayed in each panel spans from the corresponding confinement time scale t_G for the hidden-sector gauge group G, and a vertical, dashed line indicating the time scale associated with matter-radiation equality has also been included for reference. The horizontal, dash-dotted line indicates the total observed dark-matter relic abundance, as measured by the WMAP satellite. It should be emphasized that in both of these scenarios, the *total* present dark-matter relic-abundance contribution from all of the a_λ in the tower reproduces this observed value.

we will discuss the situation using the parameters which are consistent with all observational and phenomenological constraints.

Note that the sets of curves shown in the two panels of Fig. 3 differ quite significantly. For example, in the left panel, the effect of the nonzero decay widths of the heavier a_λ is readily apparent. Indeed, the curves corresponding to

masses $\lambda \gtrsim 10^{-4}$ GeV rapidly drop to zero at a time scale $t \sim \tau_\lambda$. By contrast, in the right plot, which corresponds to a scenario with $y \ll 1$, the decay rate of each a_λ is suppressed by a factor of $\tilde{\lambda}^4 A_\lambda^2$, which can be quite small in such a scenario. Consequently, all of the a_λ for which Ω_λ curves are shown in the plot are stable on cosmological time scales.
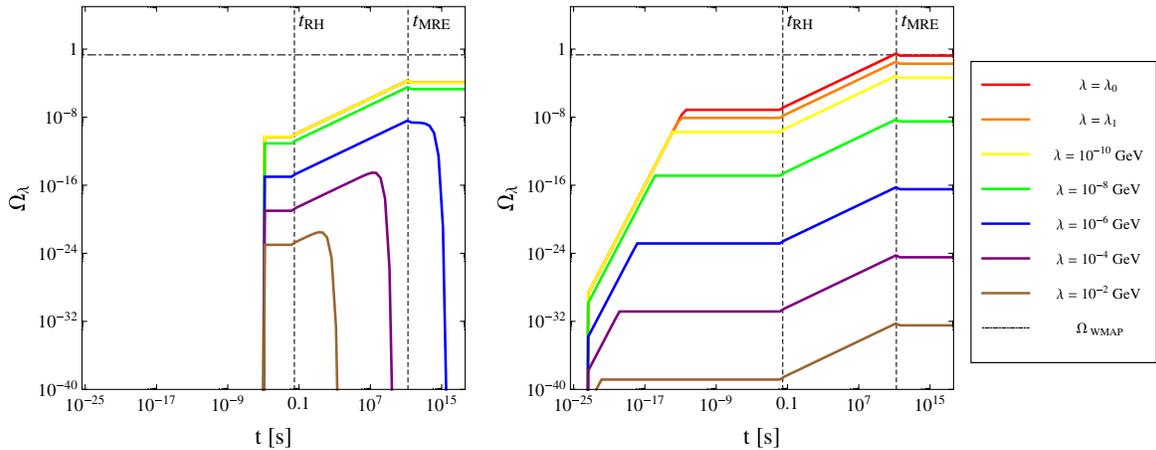
FIG. 4 (color online).   The individual relic abundances Ω_λ associated with a variety of different a_λ in the LTR cosmology, shown as a function of time t, for two scenarios with different values of $\hat{f}_X$, M_c, and Λ_G. The left panel corresponds to a choice of $\hat{f}_X = 10^6$, $M_c = 4 \times 10^{-12}$ GeV, and $\Lambda_G \approx 37$ MeV, while the right panel corresponds to a choice of $\hat{f}_X \approx 6 \times 10^{14}$ GeV, $M_c = 10^{-11}$ GeV and $\Lambda_G = 1$ TeV. In both cases, we have taken $\xi = g_G = \theta = 1$, with $T_{\rm RH} = 5$ MeV and $H_I = 10$ GeV. The range of t displayed in each panel spans from the end of cosmic inflation to present day, and vertical (dashed) lines indicating the time scales associated with the end of reheating and with matter-radiation equality have also been included for reference. The horizontal, dash-dotted line indicates the total observed dark-matter relic abundance, as measured by the WMAP satellite. It should be emphasized that in both of these scenarios, the *total* present-day dark-matter relic-abundance contribution from all of the a_λ in the tower reproduces this observed value of $\Omega_{\rm CDM}$ to within the limits quoted in Eq. (1.1).

The abundance curves shown in Fig. 4, on the other hand, display typical results obtained in the LTR cosmology. More specifically, the results shown here correspond to the parameter assignments $T_{\rm RH} = 5$ MeV and $H_I = 10$ GeV. As in Fig. 3, the two panels shown in this figure correspond to two different choices for $\hat{f}_X$, M_c, and Λ_G which both yield a total present-day relic abundance consistent with WMAP data. The curves shown in the left panel correspond to a scenario with $\hat{f}_X = 10^6$ GeV, $M_c = 4 \times 10^{-12}$ GeV, and $\Lambda_G \approx 37$ MeV. This scenario exemplifies the case in which $\hat{f}_X$ and Λ_G are both comparatively small. By contrast, the curves displayed in the right panel correspond to a scenario with far larger values for these parameters: $\hat{f}_X = 6 \times 10^{14}$ GeV and $\Lambda_G = 1$ TeV, with $M_c = 10^{11}$ GeV. In both cases, we have taken $\xi = g_G = \theta = 1$. Once again, the curves shown in each plot correspond to $\lambda = \{10^{-10}, 10^{-8}, 10^{-6}, 10^{-4}, 10^{-2}\}$ GeV, as well as λ_0 and λ_1. For the left plot, $\lambda_0 = 2 \times 10^{-12}$ GeV and $\lambda_1 = 6 \times 10^{-12}$ GeV; for the right plot, $\lambda_0 = 5 \times 10^{-12}$ GeV and $\lambda_1 = 2 \times 10^{-11}$ GeV. Again, in the left panel, the λ_0 and λ_1 curves are not apparent because they lie directly beneath the $\lambda = 10^{-10}$ GeV curve.

The contrasting features between the two panels in Fig. 4 are predominately due to the differences between their $\hat{f}_X$ and Λ_G values. First, as stated above, $\hat{f}_X$ is quite small in the scenario displayed in the left panel, and consequently, the couplings between the a_λ and the fields of the SM are quite large. This implies that the decay rates associated with the heavier a_λ will be sizable in this scenario, and that a large number of these heavier modes will decay before present time. Indeed, the precipitous drop in each of the curves corresponding to a mass eigenvalue in the range $\lambda \lesssim 10^{-6}$ GeV in this plot is a consequence of the decay of these modes to SM fields. By contrast, in the scenario displayed in the right plot, $\hat{f}_X$ is large enough that all of the a_λ for which Ω_λ curves are shown are stable on cosmological time scales, and no such effect is apparent. Second, Λ_G is also quite small in the scenario displayed in the left panel, and the confinement time scale $t_G \approx 10^{-5}$ s is consequently quite late. As a result, $t_\lambda = t_G$ for all a_λ. This implies not only that all of the modes begin oscillating at the same time, but moreover, that the Ω_λ only become nonzero quite late. By contrast, in the right panel, Λ_G is large and t_G is correspondingly quite early. This results in a situation in which the t_λ for the lighter modes are staggered in time. As discussed in Ref. [13], the primary consequence of this staggering is that the Ω_λ curves for these lighter modes depend more sensitively on λ.

The most important implication of Eqs. (4.18) through (4.21), however, is that Ω_λ decreases with increasing λ regardless of the details of the cosmological framework. On the other hand, we saw in Sec. III that Γ_λ increases monotonically with λ. This observation is indeed encouraging, in that it suggests that Γ_λ and Ω_λ possess the appropriate, reciprocal relationship needed for an ensemble of a_λ

to serve as dynamical dark matter. In the following section, we will quantify more precisely the relationship between Γ_λ and Ω_λ and demonstrate that this is indeed the case. The results of the present section therefore attest that misalignment production, in stark contrast to thermal production, is an ideal mechanism for the generation of axion relic abundances in dynamical dark-matter models. Moreover, as we shall soon demonstrate, this mechanism dominates in the regime of model-parameter space in which an ensemble of a_λ tends to be phenomenologically viable, in the sense that it correctly reproduces the observed dark-matter relic abundance, while at the same time satisfying all relevant constraints from experiment, astrophysical observation, and cosmology.

Before we proceed to analyze the collective properties of such ensembles, however, two brief comments are in order. The first of these concerns the validity of the rapid-turn-on approximation. As we have stated above, Eq. (4.6) is strictly valid only when λ and Γ_λ are essentially independent of temperature. However, there turn out to be certain situations in which the time-dependence of λ and Γ_λ at temperatures $T \gtrsim \Lambda_G$ is physically unimportant, and in which these quantities can be reliably treated as constants throughout the period in which a_λ are oscillating. One such situation arises in cases in which $t_\lambda > t_G$ for all of the a_λ which contribute meaningfully to the total dark-matter relic abundance, and therefore coherent axion oscillations do not occur until after $m_X(T)$ is effectively constant. This situation tends to arise either when Λ_G is quite large, in which case $m_X(T)$ attains its constant, late-time value very early, or else when Λ_G is fairly small, but $\hat{f}_X$ is quite large, in which case m_X itself is extremely small. As we shall see in Sec. V, these turn out to be precisely the situations in which the total relic-abundance contribution from the ensemble of a_λ successfully reproduces the observed value for $\Omega_{\rm CDM}$ quoted in Eq. (1.1). This retroactively justifies our use of the rapid-turn-on approximation.

Our second comment concerns the assumption that Γ_λ is sufficiently small that the solution of Eq. (4.6) for any given a_λ includes a period during which this equation of motion is effectively underdamped. This is critical, since in the absence of such a period, coherent oscillations cannot occur. Indeed, the energy density ρ_λ stored in any a_λ for which $\Gamma_\lambda \geq 2\lambda$ would never scale in an appropriate manner for that field to behave like massive matter; hence it would never contribute to $\Omega_{\rm CDM}$. However, it is not difficult to demonstrate that this situation essentially never arises in realistic bulk-axion scenarios, even for the most massive modes in a given tower. For example, consider the case of a purely photonic axion with $c_\gamma = 1$. In this case, it follows from Eq. (3.1) that the solution for a_λ becomes critically damped at a value $\lambda_{\rm cd}$, which is determined by the condition

$$2\lambda_{\rm cd} = G_\gamma \frac{\lambda_{\rm cd}^3}{\hat{f}_X^2} (\tilde{\lambda}^2 A_\lambda)^2. \tag{4.22}$$

Solving this equation for λ_{cd}, we find that

$$\lambda_{cd} = \frac{\hat{f}_X}{\sqrt{2G_\gamma}}\left[1 + \sqrt{1 + \frac{4G_\gamma m_X^2}{\hat{f}_X^2}\left(1 + \frac{\pi^2}{y^2}\right)}\right]^{1/2}, \quad (4.23)$$

which implies that $\lambda_{cd} \gtrsim \hat{f}_X/\sqrt{2G_\gamma}$. However, since the effective description of the theory in terms of a tower of axion modes breaks down at the five-dimensional $U(1)_X$-breaking scale $f_X \ll \hat{f}_X$, we are assured that $\Gamma_\lambda \ll 2\lambda$ for all modes in such a tower. Indeed, this qualitative result is not specific to a photonic axion, but applies broadly to any axion field which couples to the SM fields with $\mathcal{O}(1)$ coupling coefficients. As a corollary, this result also implies that the standard oscillation criterion $\lambda \sim 3H/2$ will always be met before the decay criterion $\Gamma_\lambda \sim H$. This indicates that indeed H, rather than Γ_λ, sets the time scale at which oscillations begin, and that any given a_λ decays during a time frame in which its energy density can legitimately be described by Eq. (4.13).

V. CHARACTERIZING THE ENSEMBLE: TOTAL ABUNDANCES, TOWER FRACTIONS, AND EQUATIONS OF STATE

In the previous two sections, we derived expressions for the decay widths and relic abundances for the individual mass eigenstates a_λ in a mixed tower of KK axions. We have shown that these quantities scale with λ in an appropriate, reciprocal manner for an ensemble of such states to serve as dynamical dark matter. We are therefore finally equipped to address the dark-matter phenomenology of the ensemble as a whole.

As discussed in Ref. [13], the crucial quantities which characterize a given dynamical dark-matter ensemble are the total relic abundance Ω_{tot}, the tower fraction η, and the effective equation-of-state parameter w_{eff}. In this section, we investigate how these three quantities depend on the scales $\hat{f}_X$, M_c, and Λ_G which characterize a given bulk-axion model and thereby assess which regions of parameter space are interesting from a dynamical dark-matter perspective. In the next section, we discuss the applicable phenomenological constraints on the model and demonstrate that substantial regions of parameter space exist within which all such constraints are satisfied.

A. General definitions

The first of the three principal quantities mentioned above which characterize any given dynamical dark-matter ensemble is Ω_{tot}. This is simply the total contribution to Ω_{CDM} from all constituent modes in the ensemble which have already begun oscillating:

$$\Omega_{tot} \equiv \sum_\lambda \Omega_\lambda. \quad (5.1)$$

The second is the so-called "tower fraction" η, which is a measure of how the total abundance Ω_{tot} is distributed across the ensemble. Specifically, η is defined for a given dynamical dark-matter ensemble to be the fraction of Ω_{tot} provided by all of the oscillating components of that ensemble except for the one which yields the largest individual contribution. Explicitly,

$$\eta \equiv 1 - \frac{\Omega_{max}}{\Omega_{tot}}, \quad (5.2)$$

where $\Omega_{max} \equiv \max_\lambda\{\Omega_\lambda\}$. After all the a_λ have begun oscillating, the lightest mass eigenstate a_{λ_0} always yields the largest individual relic abundance Ω_{λ_0}, and therefore $\Omega_{max} = \Omega_{\lambda_0}$. When $\eta \ll 1$, essentially the entirety of Ω_{tot} is provided by a single field, as in most traditional dark-matter models. By contrast, having $\eta \sim \mathcal{O}(1)$ signals a departure from this traditional setup, which is indeed one of the hallmarks of dynamical dark matter. (Note that when we say that η should differ significantly from zero, we are willing to accept $\eta \sim 0.1$, as such values could result in observable differences from traditional models, but not, for example, $\eta \sim 10^{-3}$.)

In our dynamical dark-matter framework, both Ω_{tot} and η are intrinsically dynamical quantities, with nontrivial time dependences. For this reason, we will designate their present-day values as $\Omega_{tot}^* \equiv \Omega_{tot}(t_{now})$ and $\eta_* \equiv \eta(t_{now})$ in what follows.

As discussed in Ref. [13], a given dynamical dark-matter ensemble as a whole can also be described in terms of a single, effective equation-of-state parameter w_{eff}:

$$w_{eff} \equiv -\left(\frac{1}{3H}\frac{d\ln\rho_{tot}}{dt} + 1\right), \quad (5.3)$$

where $\rho_{tot} \equiv \Omega_{tot}\rho_{crit}$. Indeed, one of the hallmarks of this framework is that even w_{eff} itself is continually changing in time. We shall therefore define $w_* \equiv w_{eff}(t_{now})$. As discussed in Ref. [13], this quantity is given by

$$w_* = \frac{AB}{2\Omega_{tot}^* t_{now}^{1+\alpha+\beta}} \quad (5.4)$$

for any given dynamical dark-matter ensemble in which the widths, abundances, and densities of states obey the approximate scaling relations $\Omega \approx A\Gamma^\alpha$ and $n_\Gamma \approx B\Gamma^\beta$, where n_Γ denotes the density of states per unit decay width.

Taken together, Ω_{tot}, η, and w_{eff} serve to characterize any given dynamical dark-matter ensemble. For the remainder of this section, then, our task will be to investigate how these three quantities depend on the parameters $\hat{f}_X$, M_c, and Λ_G which characterize our bulk-axion model. Also recall that in Eq. (2.21) we defined the mixing parameter

$$y = \frac{4\sqrt{2}\pi}{g_G\xi}\frac{\hat{f}_X M_c}{\Lambda_G^2}, \quad (5.5)$$

which quantifies the extent to which the different modes in the KK tower mix with each other. We will also therefore keep track of the corresponding values of y in our analysis. Moreover, as we have seen in Sect. IV, our results will also

depend on the cosmological framework adopted. We shall therefore derive results in the context of the standard and LTR cosmologies independently. However, as discussed above, constraints on T_* in theories with large, flat extra dimensions provide a strong motivation for working within the context of an LTR cosmology with a reheating temperature $T_{RH} \sim \mathcal{O}(\text{MeV})$. For this reason, the results obtained for the LTR cosmology are more likely to be realistic.

Needless to say, phenomenological consistency imposes certain constraints on the parameters Ω_{tot}, η, and w_{eff}. For example, WMAP data require that $\Omega_{tot}^* \approx \Omega_{CDM}$; likewise, w_* should not differ too significantly from zero. Beyond this, however, η and w_{eff} are fairly unconstrained by data. Nevertheless, while any values for η and w_{eff} can be realized within the general dynamical dark-matter framework, we are particularly interested in situations in which η_* is also significantly different from zero, for these are the situations in which our dynamical dark-matter ensemble represents a significant departure from traditional, single-component models of dark matter.

B. Dark towers: relic abundances and tower fractions

In Fig. 5, we show how the total present-day dark-matter relic abundance Ω_{tot}^* depends on $\hat{f}_X$, M_c, and Λ_G in the standard cosmology, assuming a photonic axion with $c_\gamma = 1$. Each panel in the figure displays contours of Ω_{tot}^* for a different choice of Λ_G. A dashed blue line highlighting the $\Omega_{tot}^* = 1$ contour has also been included in each panel. The red lines are contours of y: the solid red line corresponds to $y = 1$, which roughly indicates the transition point between the strongly-mixed regime (below and to the left of the contour) and the weakly-mixed regime (above and to the right of the contour). Proceeding from left to right, the dotted lines correspond to the values $y = \{0.01, 0.1, 10, 100\}$. In Fig. 6, we present the corresponding contour plots for the tower fraction η_*. Moreover, to complement the results shown in Figs. 5 and 6 for the standard cosmology, we present the corresponding results for the LTR cosmology in Figs. 7 and 8.

As Figs. 5 and 7 illustrate, the dependence of Ω_{tot}^* on the model parameters $\hat{f}_X$, M_c, and Λ_G is somewhat complicated, and the results displayed therein clearly warrant detailed explanation. Perhaps the most intuitive way of understanding these results is to begin by examining them in certain limiting regimes. For example, consider a situation in which Λ_G is relatively small. In this case, the confinement time scale t_G is relatively late. If t_G is sufficiently late that all modes in the tower begin oscillating immediately at t_G, each Ω_λ is given in the standard cosmology by Eq. (4.18) for all λ and by Eq. (4.19) in the LTR cosmology. In addition, let us assume that $\hat{f}_X$ is large enough that decays can be neglected, and that H_I is at least moderately large, so that essentially all of the a_λ which contribute meaningfully to Ω_{tot}^* survive inflation. In this special case, we can explicitly sum the contributions Ω_λ to obtain the result

$$\Omega_{tot}^* \approx \frac{3}{256\pi^2}(g_G\xi)^2\left(\frac{\theta\Lambda_G^2}{M_P}\right)^2 t_G^{3/2} t_{MRE}^{1/2}$$

$$\times \begin{cases} 1 & \text{standard cosmology} \\ (t_G/t_{RH})^{1/2} & \text{LTR cosmology.} \end{cases} \quad (5.6)$$

In deriving this result, we have used the second identity in Eq. (2.25). Note that Eq. (4.14) implies that $t_G < t_{RH}$; hence Ω_{tot}^* is suppressed in the LTR cosmology relative to the standard cosmology by a factor which can be quite significant. Indeed, this suppression factor in Ω_{tot}^* is due to the uniform suppression of each individual contribution Ω_λ in this regime by the factor $\mathcal{E}_{LTR}$ given for the $\lambda \geq 1/t_G$ case in Eq. (4.17).

Perhaps the most interesting aspect of this result is that it depends only on Λ_G^4, and is independent of both M_c and $\hat{f}_X$. This is quite surprising indeed, for it indicates that in this regime, no matter how many of the a_λ contribute significantly to Ω_{tot}, the total contribution to the dark-matter relic abundance is the same. For example, a strongly-mixed scenario with $y \ll 1$ and a vast number of modes contribute more or less democratically to Ω_{tot} will yield the same abundance as a weakly-mixed scenario with $y \gg 1$ in which a single, light axion accounts for essentially the entirety of the dark matter. This situation is realized in the $\Lambda_G = 10$ MeV and $\Lambda_G = 100$ MeV panels in Figs. 5 and 7, in which the value of Ω_{tot}^* remains essentially constant throughout the region of parameter space shown.

At an algebraic level, the fact that Ω_{tot}^* is a constant throughout substantial regions of $(\hat{f}_X, M_c)$ space for small Λ_G is a reflection of the fact that the identity in Eq. (2.25) holds regardless of the value of y. Of course, this identity requires that the sum over Ω_λ be taken over the entire KK tower, from the lowest mass eigenstate up to infinity. At a physical level, this is the appropriate sum to take for small Λ_G, because all of the modes begin oscillating at a common time t_G, and because the full structure of the tower is undisturbed by the decay of any modes which contribute meaningfully in the sum. Indeed, as we have seen from the panels of Figs. 5 and 7, this result is characteristic of situations in which Λ_G is small.

However, as we increase Λ_G, three effects can begin to alter this picture and thereby destroy the uniformity of Ω_{tot}^*:

(i) First, t_G becomes smaller and smaller, and consequently t_λ can begin to exceed t_G for the lower modes in the KK tower. In other words, these lower modes may begin to experience oscillations with staggered onset times, a phenomenon which begins with the lowest-lying modes in the tower and ultimately affects higher and higher modes as Λ_G increases.

(ii) Second, y decreases with increasing Λ_G (for fixed $\hat{f}_X$ and M_c), and consequently more and more of the excited a_λ contribute significantly to Ω_{tot}^*.
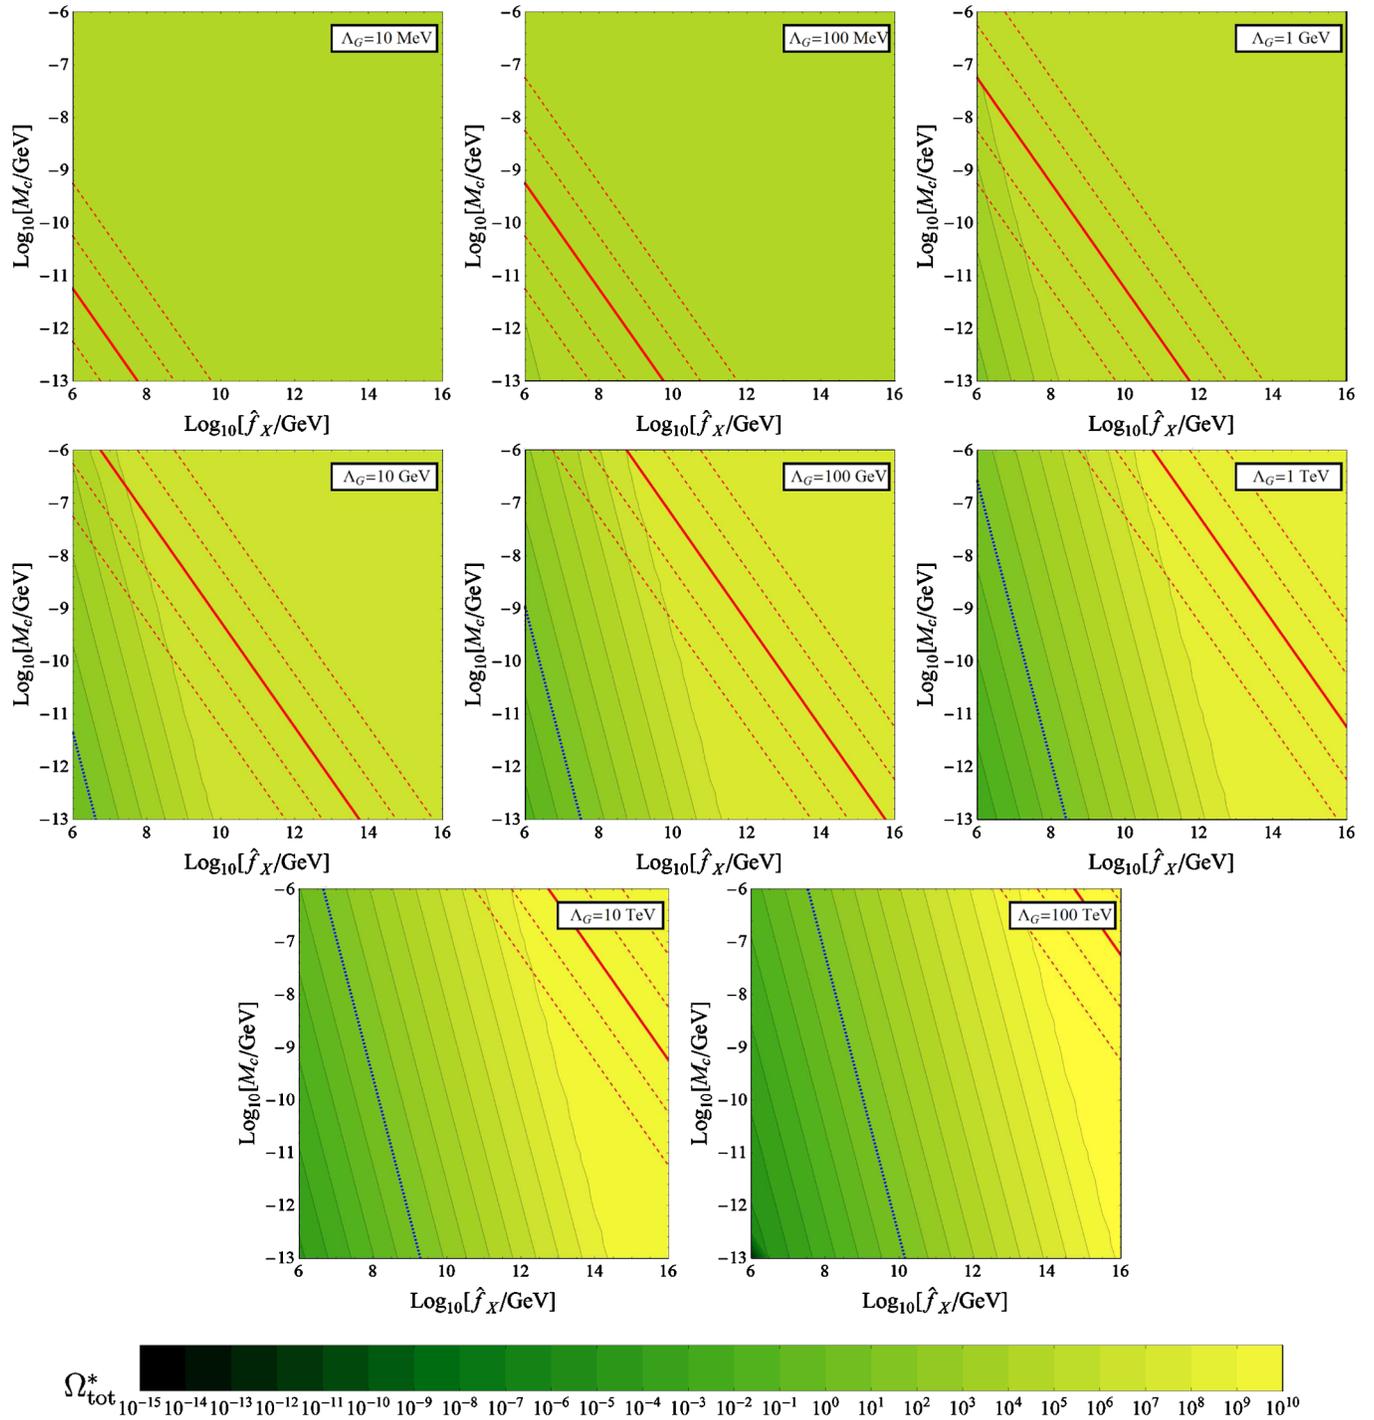
FIG. 5 (color online).   Contours of the total contribution Ω_{tot}^* from a KK tower of general axions, plotted in $(\hat{f}_X, M_c)$ space, assuming the standard cosmology. Each panel corresponds to a different choice of Λ_G ranging from 10 MeV to 100 TeV. In each case, we have taken $\xi = g_G = \theta = 1$ and set $H_I = 10^{-7}$ GeV. Note that the contour corresponding to $\Omega_{\text{tot}}^* = 1$ has been highlighted with a dotted blue line in each panel for clarity. The solid oblique red line appearing in each panel indicates where $y = 1$, and proceeding from left to right, the dashed red lines correspond to $y = \{0.01, 0.1, 10, 100\}$.

Although the lifetimes τ_λ of these modes also increase with increasing Λ_G, they do so at a slower rate. As a result, a larger and larger fraction of the contributing portion of the KK tower is effectively truncated by decays. This effect can therefore lead to a reduction in Ω_{tot}^*, especially in the $y \ll 1$ regime.

(iii) Third, as Λ_G increases, t_G can be pushed back into the inflationary era. The contributions from those modes which begin oscillating prior to or during
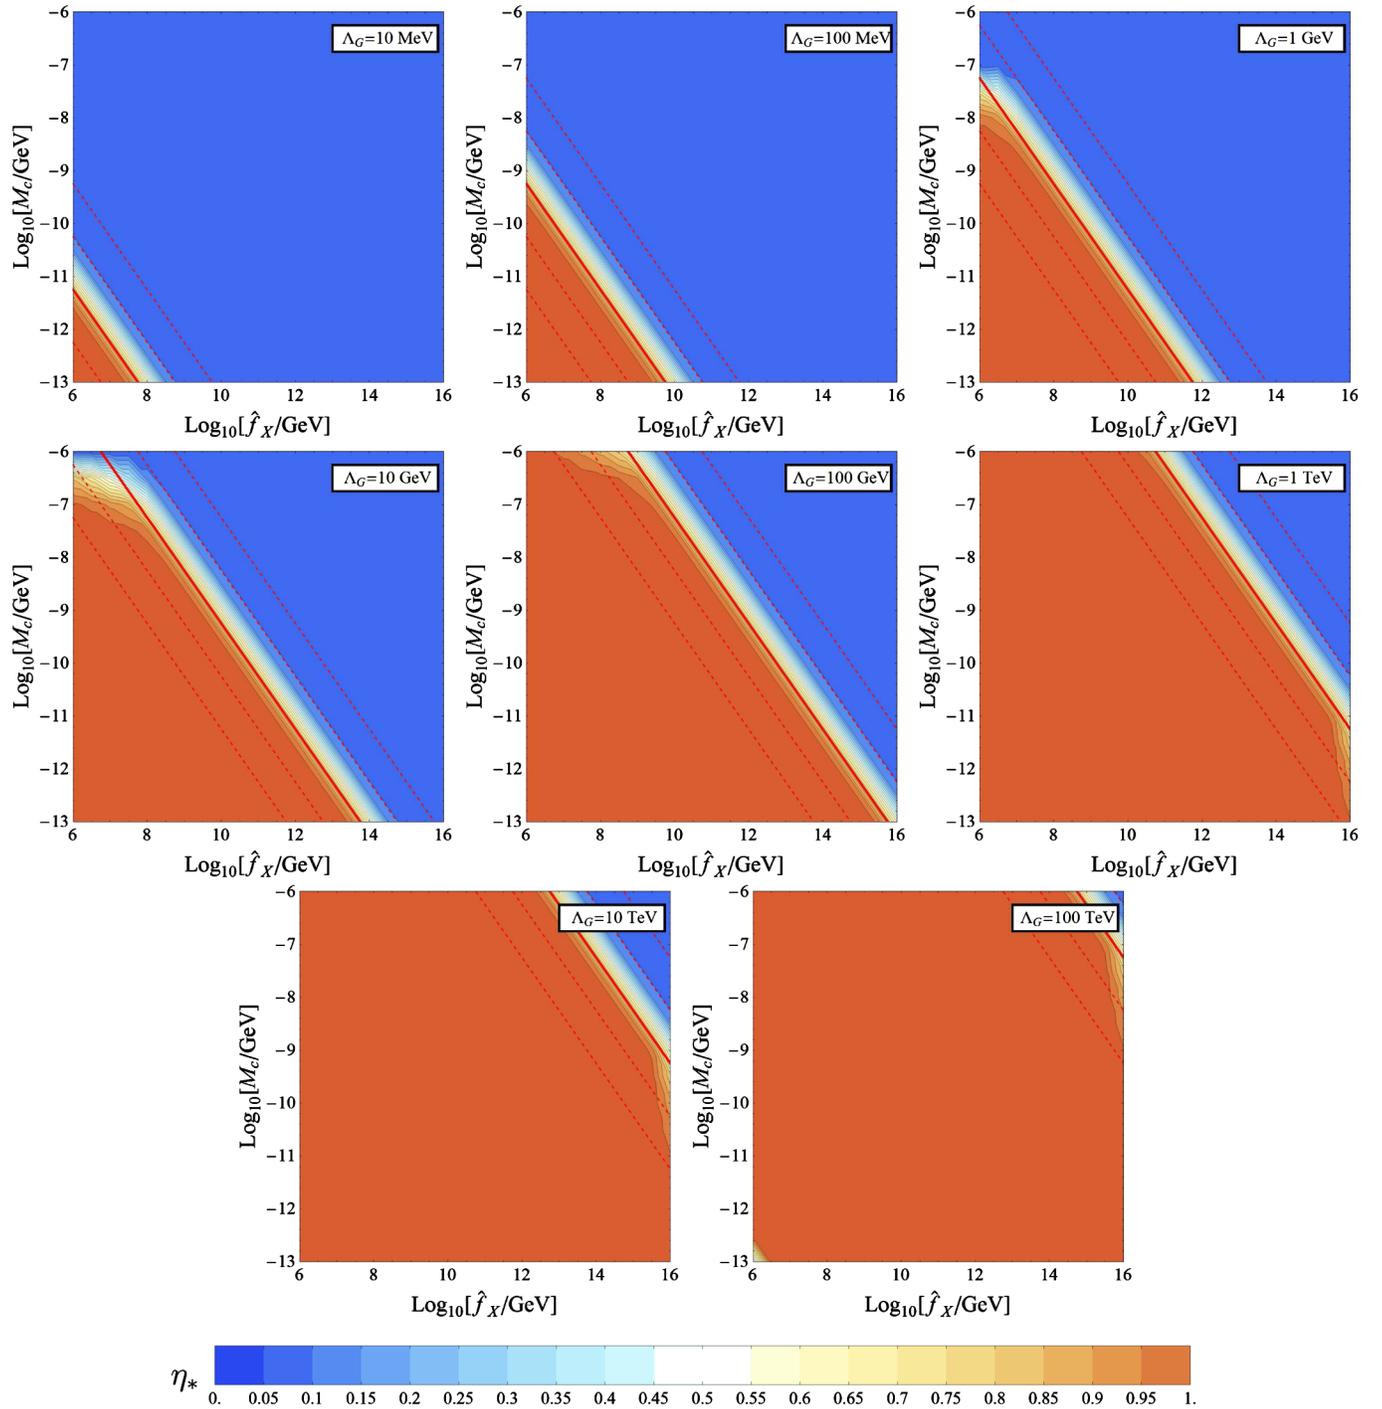
FIG. 6 (color online). Contours of the tower fraction η_* from a KK tower of general axions, plotted in $(\hat{f}_X, M_c)$ space, assuming the standard cosmology. As in Fig. 5, each panel corresponds to a different choice of Λ_G ranging from 10 MeV to 100 TeV; likewise, we have taken $\xi = g_G = \theta = 1$ and set $H_I = 10^{-7}$ GeV. Once again, the solid oblique red line appearing in each panel indicates where $y = 1$, and proceeding from left to right, the dashed red lines correspond to $y = \{0.01, 0.1, 10, 100\}$.

inflation will therefore be inflated away. This too can result in a reduction of Ω^*_{tot}.

Of course, which of these effects happens to be relevant in any given situation ultimately depends on the parameters in question, and whether we are working in the standard

cosmology or an LTR cosmology. Let us therefore begin by examining the situation in the standard cosmology, as shown in Fig. 5. As we increase Λ_G from 10 MeV to 100 TeV, we see that a series of contours with smaller and smaller values of Ω^*_{tot} emerges in the region of
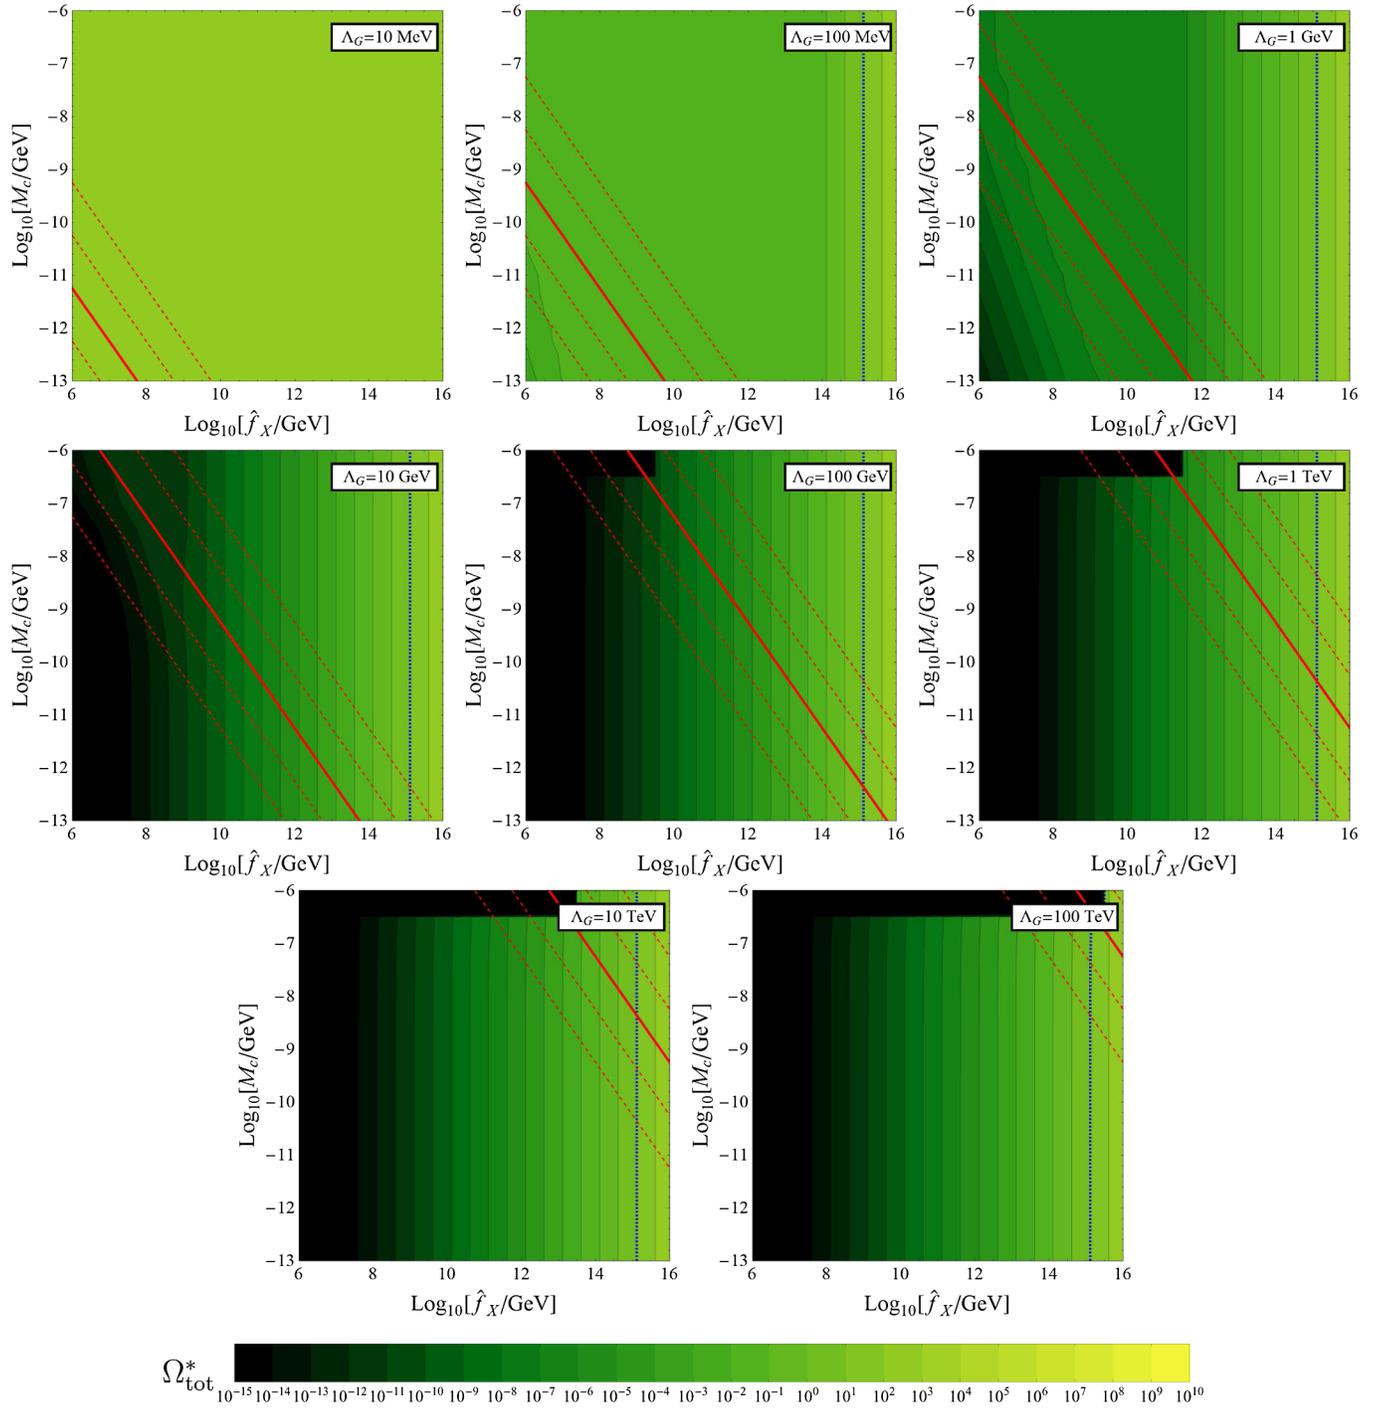
FIG. 7 (color online). Same as in Fig. 5, but for the LTR cosmology rather than the standard cosmology.

parameter space where $y \ll 1$ and gradually spreads over a substantial region of $(\hat{f}_X, M_c)$ space. This is the effect of decays truncating the contributions from the higher modes in the tower, as discussed above. Note that for the regions of parameter space shown in Fig. 5, neither of the other two effects outlined above is apparent. In particular, staggering effects only occur within regions of parameter space for which $t_G < t_{\lambda_0}$, where t_{λ_0} is the time at which the lightest mode in the tower begins oscillating. This criterion can be rephrased as a condition on the model parameters $\hat{f}_X$, M_c, and Λ_G by substituting Λ_G for T in the middle line of Eq. (4.2). In the $y \gg 1$ and $y \ll 1$ regimes, we can approximate $\lambda_0 m_X$ and $\lambda_0 \approx M_c/2$ respectively to obtain
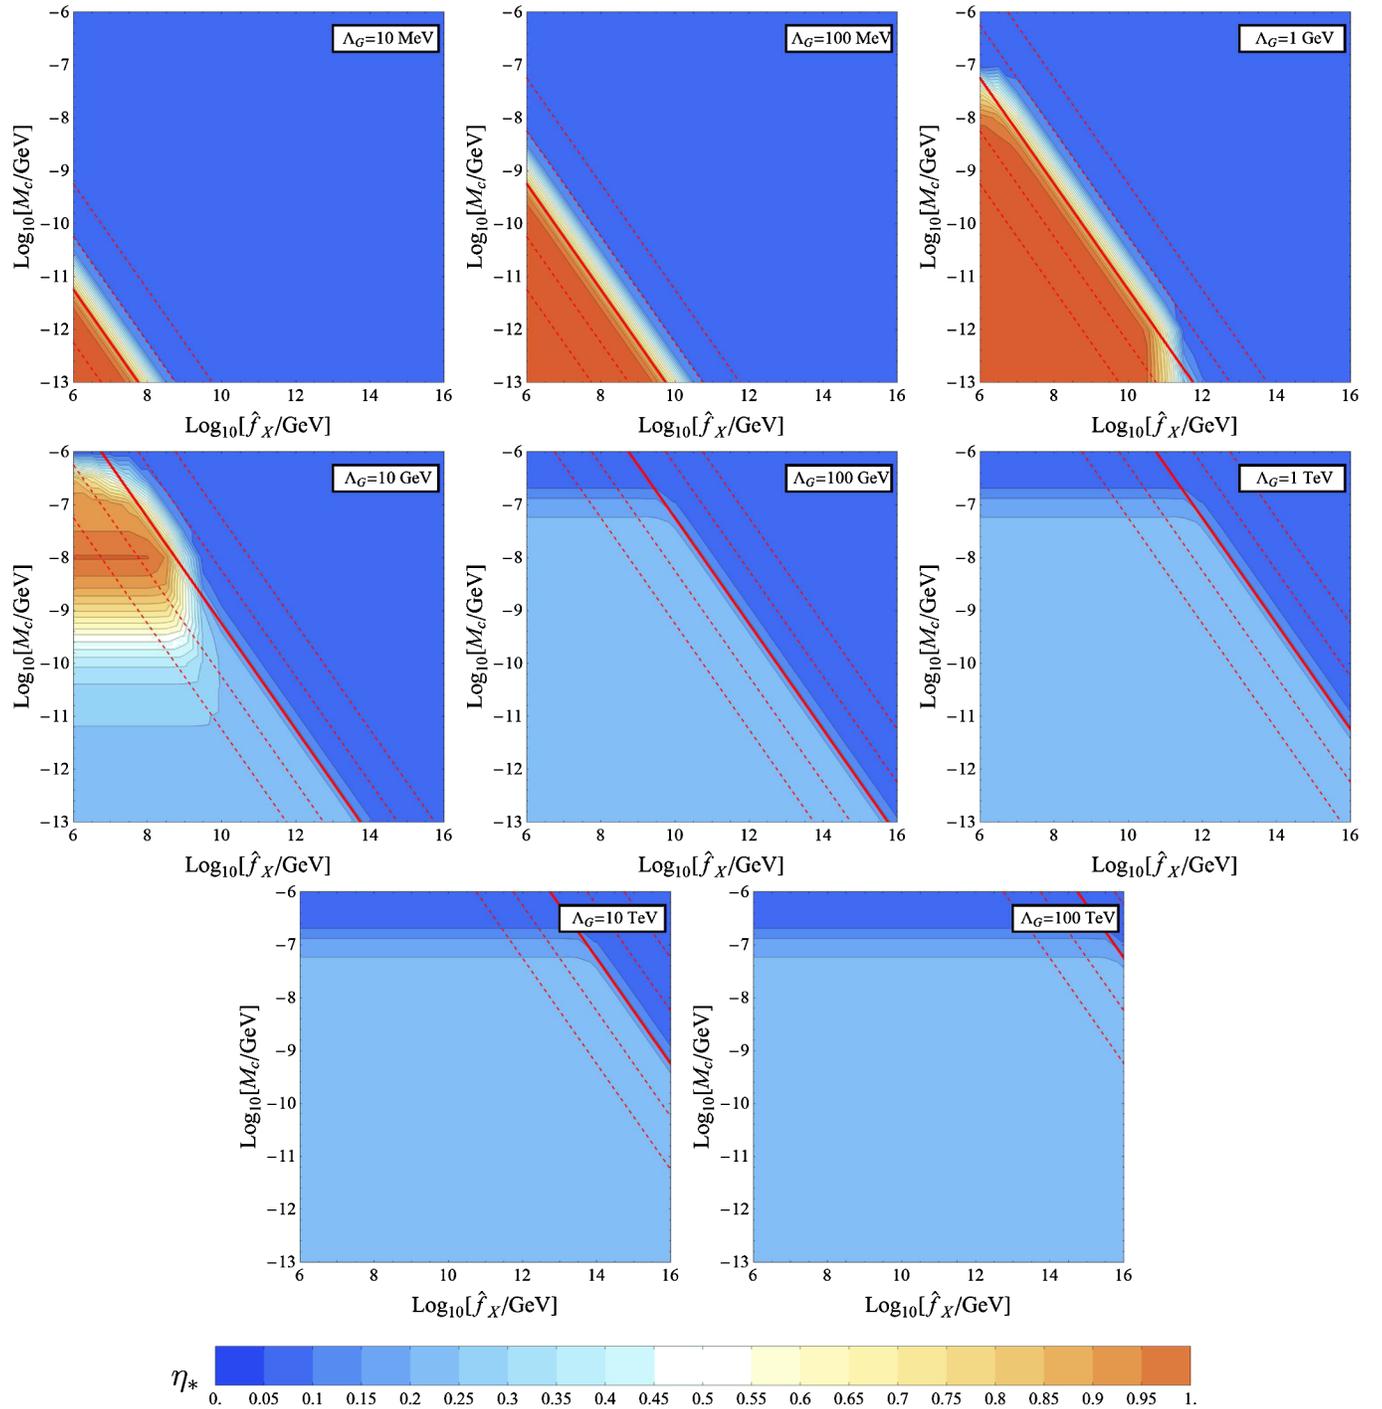
FIG. 8 (color online).   Same as in Fig. 6, but for the LTR cosmology rather than the standard cosmology.

$$\text{standard cosmology:} \begin{cases} \hat{f}_X \gtrsim (1.38 \times 10^{17} \text{ GeV}) \times g_G \xi [g_*^{-1/2}(\Lambda_G)] & y \gg 1 \\ M_c \lesssim (8.18 \times 10^{-19} \text{ GeV}) \times [g_*^{1/2}(\Lambda_G)]\left(\frac{\Lambda_G}{\text{GeV}}\right)^2 & y \ll 1. \end{cases} \quad (5.7)$$

Given these results, it is clear that staggering effects will not be visible in Fig. 5: for the $y \gg 1$ case, extremely large values of $\hat{f}_X$ are required, regardless of the value of Λ_G; for the $y \ll 1$ case, $M_c \lesssim 10^{13}$ GeV or $\Lambda_G \gtrsim 100$ TeV is required for these effects to be apparent.

The situation is quite different in the LTR cosmology, as shown in Fig. 7. Indeed, as we increase Λ_G, all three of the above effects begin to become relevant. First, we observe the same effect of decaying a_λ modes in the $y \ll 1$ region that we saw in Fig. 5. This effect is particularly evident in the $\Lambda_G = 1$ GeV panel of Fig. 7. However, in

the LTR case, we also observe effects due to staggering, which begin to appear in the large-$\hat{f}_X$ region. Indeed, following the same procedure applied above for the standard cosmology but using the top line in Eq. (4.2), we find that these effects emerge in regions of parameter space where

$$\text{LTR cosmology:} \quad \begin{cases} \hat{f}_X \gtrsim (1.03 \times 10^{11} \text{ GeV}) \times g_G \xi \left[\frac{g_*^{1/2}(T_{\text{RH}})}{g_*(\Lambda_G)} \right] \left(\frac{T_{\text{RH}}}{\text{MeV}} \right)^2 \left(\frac{\Lambda_G}{\text{GeV}} \right)^{-2} & y \gg 1 \\ M_c \lesssim (2.73 \times 10^{-13} \text{ GeV}) \times \left[\frac{g_*(\Lambda_G)}{g_*^{1/2}(T_{\text{RH}})} \right] \left(\frac{T_{\text{RH}}}{\text{MeV}} \right)^{-2} \left(\frac{\Lambda_G}{\text{GeV}} \right)^4 & y \ll 1. \end{cases} \tag{5.8}$$

The first of these limiting forms accounts for the vertical contours which appear on the right side of the $\Lambda_G = 100$ MeV panel in Fig. 7 and encroach further and further to the left as Λ_G increases.

Finally, as Λ_G grows beyond 100 GeV, we see the third effect emerging: the inflating away of heavy KK modes. In particular, since we have chosen $H_I = 10^{-7}$ GeV in this plot, the relic-abundance contributions from all modes with $\lambda \gtrsim 3H_I/2$ will be inflated away. Indeed, we see that when $y \ll 1$ (which implies that $\lambda_0 \approx M_c/2$) and $M_c \gtrsim 3 \times 10^{-7}$ GeV, the entire tower is inflated away, yielding $\Omega_{\text{tot}}^* = 0$. Indeed, since y increases with Λ_G for fixed $\hat{f}_X$

and M_c, this effect spreads across a wider region as Λ_G increases.

Ultimately, for large Λ_G, those modes which have not inflated away exhibit a completely staggered behavior in the LTR cosmology. This limit may be regarded as the converse of the "instantaneous turn-on" limit taken in Eq. (5.6) for small Λ_G: indeed, we now have $t_\lambda > t_G$ for *all* of the modes which contribute significantly to Ω_{tot}^*. Moreover, in this case the sum over Ω_λ can be explicitly evaluated using the first identity in Eq. (2.25), allowing us to obtain explicit results for Ω_{tot}^* in the completely staggered limit for both the standard and LTR cosmologies:

$$\text{standard cosmology:} \quad \Omega_{\text{tot}}^* \approx \frac{3^{5/4}}{2^{29/4}\pi^{1/2}} (g_G\xi)^{1/2} \left(\frac{\theta}{M_P} \right)^2 t_{\text{MRE}}^{1/2} \hat{f}_X^{3/2} \Lambda_G C(y)$$

$$\text{LTR cosmology:} \quad \Omega_{\text{tot}}^* \approx \frac{3}{8} \left(\frac{\theta}{M_P} \right)^2 \left(\frac{t_{\text{MRE}}}{t_{\text{RH}}} \right)^{1/2} \hat{f}_X^2, \tag{5.9}$$

where $C(y) \equiv \sum_\lambda \tilde{\lambda}^{1/2} A_\lambda^2$ for the standard-cosmology case. Note that the y-dependence of $C(y)$ is illustrated in Fig. 4 of Ref. [13]. In sharp contrast with the results obtained in Eq. (5.6), we see that in this staggered regime, the expressions for Ω_{tot}^* in the standard and LTR cosmologies differ significantly. In the standard cosmology, Ω_{tot}^* depends nontrivially on $\hat{f}_X$, Λ_G, and M_c (through its dependence on y). By contrast, in the LTR cosmology, Ω_{tot}^* depends on $\hat{f}_X$ in this staggered-oscillation regime, but not on Λ_G or M_c.

Following similar reasoning, we can also understand the behavior of the tower fraction η_*, beginning with the results displayed in Fig. 6 for the standard cosmology. These results nicely illustrate an important general property of η_* in brane/bulk models of dynamical dark matter discussed in Ref. [13], which is that the behavior of η_* is strongly correlated with the value of y. In particular, when $y \gg 1$, the mass of the lightest mode in the KK tower is proportionally far lighter than those all of the excited modes, and consequently its contribution Ω_{λ_0} to the total abundance will be much larger than the contributions from all of those other modes combined. Indeed, this is nothing but the four-dimensional limit of the KK theory. This property of η_* is independent of both the specific

cosmological context and whether or not any of the modes in the tower have staggered oscillation onset times, as can be seen from Eqs. (4.18) through (4.21). By contrast, when $y \ll 1$, each mode with a mass $\lambda \ll \lambda_{\text{trans}}$ contributes essentially equally toward Ω_{tot}^* when $t_\lambda = t_G$ for all a_λ, and hence $\eta_* \approx 1$. This behavior is manifest in the various panels of Fig. 6, in which η_* rapidly transitions from nearly zero to nearly unity as one crosses the $y = 1$ contour.

In Fig. 8, we display the behavior of η_* in the LTR cosmology. For small Λ_G, the situation is very similar to that in the standard cosmology: all modes in the tower begin oscillating at t_G, and hence $\eta \approx 0$ for $y \gg 1$, while $\eta \approx 1$ for $y \ll 1$. However, as Λ_G increases, the oscillation onset times for more and more of the lighter modes in the tower become staggered in time. In this regime, the energy densities ρ_λ associated with the lighter modes in the tower scale like vacuum energy (and are hence unaffected by Hubble dilution) for a longer time before coherent oscillations set in and they begin to scale like massive matter, as illustrated in Fig. 1 of Ref. [13]. It follows that in this regime, the lighter modes account for a greater fraction of Ω_{tot}^*. For this reason, as discussed in Ref. [13], η_* no

longer approaches unity for $y \ll 1$ in those regions of parameter space in which oscillation onset times are staggered, but instead asymptotes to[1]

$$\eta_{\max} \equiv 1 - \frac{8}{\pi^2} \approx 0.189, \qquad (5.10)$$

for the case of the LTR cosmology. Roughly speaking, the regions of parameter space in which this occurs are those in which the staggered-onset criteria in Eq. (5.8) are satisfied. Indeed, this effect first becomes apparent in the $\Lambda_G = 1$ GeV panel of Fig. 8 and becomes increasingly significant as Λ_G increases and staggering effects become relevant for smaller and smaller $\hat{f}_X$ and larger and larger M_c. By $\Lambda_G = 10$ GeV, these staggering effects are realized over nearly the entirety of $(\hat{f}_X, M_c)$ space shown, leaving only a narrow strip in which all modes still begin oscillating at t_G, and by $\Lambda_G = 100$ GeV, even this strip vanishes. Note also that the effect on η_* of modes being inflated away is apparent in the upper left of those panels in Fig. 8 for which $\Lambda_G \gtrsim 10$ GeV. While η is technically undefined in this region of parameter space because $\Omega_{\mathrm{tot}}^* = 0$, we have set $\eta_* = 0$ within this region to illustrate where this effect is relevant.

As discussed in the beginning of this section, the interesting regions of parameter space for dynamical dark matter are ultimately those in which $\Omega_{\mathrm{tot}}^* \approx \Omega_{\mathrm{CDM}}$, while at the same time η_* differs significantly from zero. Given the results in Figs. 5 through 8, we can now determine whether this situation ever actually arises in our model. Comparing the results in Figs. 5 and 6, we see that this occurs in the standard cosmology for small values of $\hat{f}_X$, within a diagonal stripe of parameter space slightly to the left of the blue $\Omega_{\mathrm{tot}}^* = 1$ contour in each panel. We also see that this stripe moves to the right in $(\hat{f}_X, M_c)$ space as Λ_G increases. By contrast, comparing the results in Figs. 7 and 8, we see that the above conditions are satisfied in the LTR cosmology in the region of parameter space where

$$\text{preferred region (LTR):} \begin{cases} \bullet \ \hat{f}_X \sim 10^{14}\text{–}10^{15} \text{ GeV} \\ \bullet \ \Lambda_G \gtrsim 100 \text{ GeV} \\ \bullet \ M_c \text{ small enough that } y \lesssim 1. \end{cases}$$

$$(5.11)$$

This result is certainly intriguing, as it suggests that the preferred scale for Λ_G in this model is roughly the TeV scale for the LTR cosmology—a scale at which there is good reason to expect new physics to appear.

The principal message of Figs. 5 through 8, then, is that our bulk-axion model indeed satisfies the conditions on Ω_{tot}^* and η_* for dynamical dark matter within these regions of parameter space. In other words, within these regions, our axion ensemble reproduces the observed dark-matter relic abundance, and does so in a nontrivial manner, with a substantial number of its constituents contributing significantly to Ω_{CDM}. Of course these alone are not sufficient conditions for a successful model of dynamical dark matter: such a model must also not only have an appropriate present-day equation-of-state parameter w_*, but also satisfy all additional relevant phenomenological constraints. In the remainder of this section, we will address the constraints on w_*; the rest of the applicable constraints will be addressed in Sect. VI.

One particular ramification of these constraints, however, is appropriate to mention before proceeding further. As discussed in Sec. IV, certain bounds which apply generically to models with large, flat extra dimensions strongly prefer the LTR cosmology over the standard cosmology. For this reason, we will focus primarily on the LTR case from this point forward.

C. Dark towers: equations of state

Having characterized the behavior of Ω_{tot}^* and η_* over the parameter space of our bulk-axion model, we now proceed to discuss the third critical quantity which characterizes the dynamical dark-matter ensemble in this model: the present-day effective equation-of-state parameter w_*. Since it is now clear which regions of model parameter space are suitable for dynamical dark matter, we will not perform a general survey of w_* over the entirety of that parameter space, as we did with Ω_{tot}^* and η_*, but instead focus on the preferred regions indicated in Eq. (5.11).

In order to calculate w_* we need to know the values of the coefficients and exponents A, B, α, and β appearing in Eq. (5.4). This, in turn, requires knowledge of how our abundances and decay widths scale with λ. As in the previous subsection, we will assume that the abundances of the a_λ result from misalignment production, and likewise we will assume that their decay widths Γ_λ are those appropriate for a photonic axion with $c_\gamma = 1$. Because the preferred region of parameter space specified in Eq. (5.11) for our model is one which is well approximated by assuming staggered oscillation onset times for all relevant modes, the correct expression for Ω_λ is the one given in Eq. (4.21). Likewise, the decay width for a photonic axion is given by the expression in Eq. (3.1). We therefore find that the coefficients A and B appearing in Eq. (5.4) are given, respectively, by

---

[1]Note that there are two (ultimately equivalent) ways to derive this result, corresponding to two different methods of taking the $y \to 0$ limit. In Ref. [13], we recognized that $\lambda_n \approx (n + 1/2)M_c$ as $y \to 0$. Since the $y \to 0$ limit also implies that $\lambda \ll \lambda_{\mathrm{trans}}$ for all λ, we can similarly approximate $A_\lambda \sim 1/\lambda$ in this limit. We then have $\Omega_\lambda \sim A_\lambda^2 \sim 1/\lambda^2$, whereupon it follows that $\Omega_{\lambda_0}/\Omega_{\mathrm{tot}} = 4/\sum_n (n + 1/2)^{-2} = 8/\pi^2$, or equivalently $\eta_{\max} = 1 - 8/\pi^2$. However, it is also possible to retain the exact form $\Omega_\lambda \sim A_\lambda^2$, whereupon we see that $\Omega_{\lambda_0}/\Omega_{\mathrm{tot}} = A_{\lambda_0}^2$ where A_{λ_0} is the value of A_λ for the lightest eigenvalue λ_0 and where we have used the identity $\sum_\lambda A_\lambda^2 = 1$ to perform the sum over KK modes. Note that this result is exact and valid for all y. However, it is easy to verify that $A_{\lambda_0} \to 2\sqrt{2}/\pi$ as $y \to 0$. We thus again find that $\eta_{\max} = 1 - 8/\pi^2$.

$$A = 3\frac{\theta^2}{M_P^2}\left(\frac{t_{\rm MRE}}{t_{\rm RH}}\right)^2 \times \begin{cases} \frac{2^{4/3}G_\gamma^{4/3}m_X^4}{\hat{f}_X^{2/3}} & \lambda \gtrsim \frac{\pi m_X^2}{M_c} \\ \frac{2^{2/5}G_\gamma^{4/3}(\hat{f}_X m_X)^{6/5}}{(1+\pi^2/y^2)^{7/5}} & \lambda \lesssim \frac{\pi m_X^2}{M_c} \end{cases} \quad (5.12)$$

and

$$B = \begin{cases} \frac{(2\hat{f}_X m_X)^{2/3}}{6M_c G_\gamma^{1/3}} & \lambda \gtrsim \frac{\pi m_X^2}{M_c} \\ \frac{(4\hat{f}_X m_X)^{2/5}}{10M_c G_\gamma^{1/5}}(1+\pi^2/y^2)^{1/5} & \lambda \lesssim \frac{\pi m_X^2}{M_c}, \end{cases} \quad (5.13)$$

while the power-law indices α and β are given by

$$(\alpha,\beta) \approx \begin{cases} (-4/3,-2/3) & \lambda \gtrsim \frac{\pi m_X^2}{M_c} \\ (-2/5,-4/5) & \lambda \lesssim \frac{\pi m_X^2}{M_c}. \end{cases} \quad (5.14)$$

Substituting these results into Eq. (5.4), we find that

$$w_* = \frac{\theta^2}{M_P^2 M_c \Omega_{\rm tot}^*}\left(\frac{t_{\rm MRE}}{t_{\rm RH}}\right)^{1/2} \times \begin{cases} G_\gamma m_X^4 t_{\rm now} & \lambda \gtrsim \frac{\pi m_X^2}{M_c} \\ \frac{3(2G_\gamma m_X^8 \hat{f}_X^8 t_{\rm now})^{1/5}}{10(1+\pi^2/y^2)^{6/5}} & \lambda \lesssim \frac{\pi m_X^2}{M_c}. \end{cases}$$
$$(5.15)$$

Let us discuss the implications of these results. First, it was noted in Ref. [13] that the effective equation-of-state parameter $w_{\rm eff}(t)$ for any given dynamical dark-matter ensemble at any time $t < t_{\rm now}$ will always fall within the range $0 \leq w_{\rm eff}(t) \leq w_*$ as long as $\alpha + \beta < -1$. This makes such ensembles less dangerous from a phenomenological point of view. Indeed, we see from the results above that this criterion is satisfied for both the large-λ and small-λ regimes in the bulk-axion model under consideration here.

Second, in order to convey a sense of the characteristic size of w_* in the favored region of parameter space for dynamical dark matter given in Eq. (5.11), we note that for the choice of $\hat{f}_X = 10^{14}$ GeV, $M_c = 10^{-11}$ GeV, and $\Lambda_G = 1$ TeV, with $g_G = \xi = \theta = 1$, we find that $w_* \approx 8.4 \times 10^{-23}$ for $\lambda \gtrsim \pi m_X^2/M_c$, while $w_* \approx 5.7 \times 10^{-11}$ for $\lambda \lesssim \pi m_X^2/M_c$. As these numbers are both extremely close to zero, we conclude that at present time our axion ensemble has an effective equation of state which can be legitimately interpreted as that of dark matter. Thus our ensemble meets all three requirements for a self-consistent model of dynamical dark matter.

VI. CHARACTERIZING THE ENSEMBLE: CONSTRAINTS AND PROSPECTS FOR DETECTION

In the previous section, we demonstrated that an ensemble of mixed KK excitations of a bulk-axion field can collectively account for the observed relic abundance of dark matter in our universe. However, as discussed in Ref. [13], in order to be a viable model of dynamical

dark matter, the model must also comply with a variety of additional laboratory, astrophysical, and cosmological constraints. Some of these constraints are intrinsic to any theory involving large extra dimensions, while others arise due to the physical effects of the axion field which propagates in the bulk of those dimensions. A number of analyses of such constraints exist in the literature [26,41,42] for the specific case in which the bulk axion in question is identified with the QCD axion and the fundamental, D-dimensional quantum-gravity scale is taken to be roughly $M_D \sim \mathcal{O}({\rm TeV})$. By contrast, in the present analysis, we are interested in a broader class of axions which are neither required to couple to the fields of the SM (and, in particular, to hadrons) in the same manner as a QCD axion, nor subject to the same strict relationship between the suppression scale for those couplings and the axion mass. Moreover, our primary motivation is not to address the hierarchy problem, but to address the issue of what constitutes the nonbaryonic dark matter in our universe. For these reasons, we will not focus exclusively on scenarios in which M_D is at or near the TeV scale, but also consider scenarios with much larger M_D. As a consequence, exclusion limits on the parameter space of the more general axion scenarios considered here can differ quite significantly from those presented in previous studies, and thus warrant reexamination.

We begin our summary of the applicable constraints on our model with a brief synopsis of those limits which arise generically in theories with large, flat extra dimensions and which do not depend on the presence or properties of the bulk-axion field. For the most part, these limits, an overview of which was presented in Ref. [28], tend to derive from the nonobservation of physical effects related to the dynamics of KK gravitons. These limits take many forms. First, there is the direct lower bound on M_c quoted in Eq. (3.4) from experimental limits on modifications of Newton's law at short distances due to KK-graviton exchange [33]. In addition, a number of constraints arise as a consequence of the production of these particles in the early universe [28,36]. As discussed in Sec. IV, these cosmological constraints can collectively be addressed by positing that the universe underwent a late period of cosmic inflation with a reheating temperature $T_{\rm RH} \sim \mathcal{O}({\rm MeV})$. Thus, by adopting an LTR cosmology with a reheating temperature of this order, as we have done, we automatically ensure that a large number of these model-independent constraints are satisfied.

A number of additional constraints on theories of this sort can be derived from observational limits on KK-graviton production in astrophysical sources, such as stars [43] and supernovae [44,45]. The most stringent of these constraints are currently those resulting from gravitationally trapped KK gravitons in the halos of neutron stars either decaying to photons or serving as a heat source for the stars themselves. In the case of $n > 1$ flat

extra dimensions with equal radii compactified on an
n-torus, these limits supersede the limit on M_c given
in Eq. (3.4). In particular, for $n = 2$, the bound is $M_c \gtrsim$
5.8×10^{-7} GeV, while for $n = 3$, one finds $M_c \gtrsim$
3.8×10^{-10} GeV [43]. However, if the radii of the extra
dimensions differ from one another, or if the compactifi-
cation manifold is not toroidal, these bounds can be con-
siderably weaker. Furthermore, it is possible that the axion
propagates only within some number n_a of the additional
dimensions, $n_a < n$. In other words, the axion could be
confined to a $(4 + n_a)$-dimensional brane within the bulk.
In this case, the effective, four-dimensional scales $\hat{f}_X$ and
M_P are related to the fundamental, higher-dimensional
scales f_X and M_D in completely different ways:

$$M_P^2 = V_n M_D^{2+n} \qquad \hat{f}_X^2 = V_{n_a} f_X^{2+n_a}. \qquad (6.1)$$

The upshot, then, is that naïve limits on M_D derived from
KK-graviton dynamics under the assumption of toroidal
compactification and equal radii do not necessarily trans-
late in a straightforward manner into constraints on the
mass scales relevant to the physics of a bulk axion.
Fortunately, the bound in Eq. (3.4) is universal and is not
sensitive to the total number of extra dimensions, unless
they are each of comparable size. We will therefore take
this bound to be the lower limit on M_c in the $n_a = 1$ model
under consideration here.

We now turn to address those constraints which relate to
the effects of the bulk axion itself. Indeed, a number of
considerations serve to constrain the properties of light
exotic particles with suppressed couplings to SM fields.
Some of these constraints derive from observational limits
on the production of such particles in astrophysical sources
such as stars and supernovae; others derive from limits on
the decays of a cosmological population of such fields into
SM fields; and still others owe to direct experimental
bounds from microwave-cavity experiments, helioscopes,
etc. A detailed analysis of the exclusion limits implied by
these constraints on general bulk-axion scenarios will be
presented in Ref. [32]. Here, we merely summarize the
results and discuss their implications for a mixed KK tower
of axions as a model of dynamical dark matter.

As we shall discuss further in Ref. [32], it is convenient
to separate the applicable constraints into four rough
classes, based on the origin of the constraint and on the
dynamics being probed. The first class of constraints which
apply to scenarios of this sort are those related to the total
present-day dark-matter relic abundance Ω_{tot}^*. Most of
these bounds have been addressed in previous sections,
but it will be useful to recapitulate them here:

(i) The axion ensemble must yield an acceptable con-
tribution to the present-day dark-matter relic density.
While $\Omega_{tot}^* < \Omega_{CDM}$ is permitted, provided some
additional field or fields make up the deficit, values
of Ω_{tot}^* in excess of the WMAP upper bound in
Eq. (1.1) are excluded.

(ii) At no time in the past may our ensemble overclose
or prematurely matter-dominate the universe.

(iii) The present-day effective equation-of-state pa-
rameter w_* for the ensemble must not deviate
significantly from zero.

(iv) Misalignment production must provide the domi-
nant contribution to Ω_λ for all a_λ, and the popula-
tion of hot axions generated via thermal production
must be negligible. We therefore require that
$\Gamma_{prod} \ll H$ at all times after the end of cosmic
inflation, where Γ_{prod} is the total production rate
of axions from interactions with SM fields in the
thermal bath.

(v) We have also assumed that the population of axions
generated from the decays of cosmic strings associ-
ated with the breaking of the global $U(1)_X$ symme-
try is small compared to the population generated by
misalignment production. We therefore impose the
requirement that $f_X \gtrsim H_I$, so that such strings are
diluted away by inflation.

(vi) Our model must respect current observational lim-
its on isocurvature fluctuations from WMAP [2].

The last of these constraints warrants additional discus-
sion. Nonadiabatic fluctuations—also known as isocurva-
ture fluctuations—refer to fluctuations not in the total
energy density (which relates directly to spacetime curva-
ture) but rather in how that total energy density is distrib-
uted among different contributing fields (including the
collective contribution from the dark sector). Such isocur-
vature fluctuations are tightly constrained by a combina-
tion of CMB observations, baryon-acoustic-oscillation
(BAO) measurements, and supernova data [2]. Such fluc-
tuations generically arise whenever a cosmological popu-
lation of particles is produced in a manner such that its
primordial density perturbations are uncorrelated with
those of the inflaton field. Indeed, limits on isocurvature
fluctuations place severe constraints on the relic abundance
of a standard QCD axion produced via vacuum misalign-
ment [17], so it might reasonably be assumed that such
limits might play a significant role in constraining our
model as well.

It turns out, however, that our model satisfies the WMAP
constraints on nonadiabatic fluctuations far more easily
than do standard axion dark-matter models. A detailed
discussion of these constraints and how they apply to
bulk-axion models of dynamical dark-matter will be pre-
sented in Ref. [32], but the gist of the argument is as
follows. Although our dynamical dark-matter ensemble
comprises a large number of individual components a_λ,
the fact that $\Omega_{tot}^* \approx \Omega_{CDM}$ implies that the individual
abundance Ω_λ associated with each of these components
is actually quite small. Furthermore, the underlying five-
dimensional nature of our KK axion tower guarantees that
the primordial density fluctuations for each a_λ are all
determined by the fluctuations $\delta\theta$ of the *same* initial

misalignment angle θ. For these reasons, the expected magnitude for isocurvature fluctuations in our model turns out to be no greater than it is in models in which misalignment production causes a single four-dimensional field to carry the complete dark-matter abundance. Moreover, if one assumes a Gaussian distribution for $\delta\theta$, it is straightforward to demonstrate [32] that $\langle(\delta\theta)^2\rangle \sim H_I^2/(2\pi\hat{f}_X)^2$. Thus, all that is required is that $H_I \ll \hat{f}_X$ within our preferred regions of parameter space. However, as discussed in Eq. (5.11), the phenomenologically preferred scale for $\hat{f}_X$ in our model is roughly $\mathcal{O}(10^{14}$–$10^{16})$GeV. For $\hat{f}_X$ at or around this scale, it turns out that current constraints on isocurvature fluctuations can be satisfied, provided that $H_I \lesssim \mathcal{O}(10^9$–$10^{10})$ GeV. Such a scale for H_I is easy to realize in traditional cosmological scenarios, and is even more natural in LTR cosmologies wherein the reheating temperature is $\mathcal{O}(\text{MeV})$. Thus, in our model, it is not difficult to satisfy current isocurvature bounds while simultaneously obtaining a total relic abundance $\Omega_{\text{tot}}^* \approx \Omega_{\text{CDM}}$.

The underlying reason why our model easily evades these nonadiabatic constraints is that within the preferred region of parameter space in Eq. (5.11), the five-dimensional axion in our model is not the standard QCD axion. In particular, we see that the scale Λ_G is significantly larger than Λ_{QCD}. Our model is thus freed from the implicit parametric dependence on Λ_{QCD} which afflicts more traditional models of axion dark matter, and allows the corresponding nonadiabatic fluctuations to have a much smaller scale.

A second class of constraints comprises those observational limits on processes in which axions are produced via their interactions with the fields of the SM and then subsequently detected via those same interactions. These include:

 (i) Limits from helioscope experiments, such as CAST [46], which search for axions produced by interactions with SM particles in the sun via their "conversion" to photons in the presence of a magnetic field.
 (ii) Limits from light-shining-through-walls (LSW) experiments (see Ref. [47] for a thorough review), including those by the BEV and GammaeV collaborations.

The most stringent of these bounds is currently that from CAST; we shall therefore take the CAST bound as representative of this class.

The physical processes to which this second class of limits applies are all subject to a particular effect which arises universally in models with both brane and bulk mass terms. This is the phenomenon of decoherence discussed in Refs. [13,26]. This decoherence phenomenon can substantially suppress the cross sections for such processes in our model, and thereby significantly weaken the bounds on $\hat{f}_X$, M_c, and Λ_G. To summarize, the cross section for any

process in which axions are produced at some time t_0 and then subsequently detected at a later time t is given by

$$\sigma(t) \propto \frac{N^2}{\hat{f}_X^4} P(t), \tag{6.2}$$

where $N \sim f_X/M_c$ is the number of modes contributing in the sum and where $P(t)$ is the detection probability at time t. This latter quantity is given in the relativistic limit by [26]

$$P(t) = \frac{1}{N^2}\left[\sum_\lambda \tilde{\lambda}^8 A_\lambda^4 e^{-\Gamma_\lambda t} + \sum_\lambda \sum_{\lambda'\neq\lambda} \tilde{\lambda}^4 \tilde{\lambda}'^4 A_\lambda^2 A_{\lambda'}^2 \right.$$
$$\left. \times e^{-(\Gamma_\lambda+\Gamma_{\lambda'})t/2} \cos\left(\frac{(\lambda^2-\lambda'^2)(t-t_0)}{2p}\right)\right], \tag{6.3}$$

where p is the initial momentum of the axion. For any reasonable choice of model parameters, the sum in the second term decoheres on time scale so rapid as to be effectively instantaneous [26]. As a result, $\sigma(t)$ is suppressed, relative to the naïve expectation, by an additional factor of N. This effect considerably weakens the constraints in this class.

A third class of constraints can be derived from processes in which axions are produced via their interactions with SM fields but not subsequently detected. Instead, the presence of the axions is made manifest by their ability to carry away momentum and energy from a given system. These constraints include:

 (i) Observational limits on the energy loss in supernovae, and, in particular, on the fraction of the energy released by SN1987A in the form of light exotic fields [48].
 (ii) Limits related to the effects of energy dissipation by axions on stellar lifetimes. The most stringent such limits currently come from observations of globular-cluster stars [33], but similar limits have also been derived from constraints on the lifetimes or energy-loss rates of other astrophysical bodies (e.g., the sun [49] and white dwarfs [50]).
 (iii) Constraints from the absence of observed signals in channels such as $j + \not{E}_T$ and $\gamma + \not{E}_T$ at particle colliders. In general, the constraints on axion production in these channels are analogous to the well-known constraints on KK-graviton production [51].
 (iv) Limits on the branching fractions in particular exotic decay channels for certain hadrons [52].

The degree to which many of these limits constrain the parameter space of bulk-axion scenarios depends quite crucially on how the axion in question couples to the fields of the SM. Moreover, many of the constraints in this class are considerably relaxed in regions of parameter space in which $y \lesssim 1$, due to the coupling-suppression phenomenon discussed in Ref. [13]. This effect will be discussed in greater detail in Ref. [32].

A fourth and final class of constraints is related to the interactions and decays of a cosmological population of axions. Depending on the cosmological epoch during which such decays occur, they can result in a number of potential signals, none of which have been observed to date. For example, these include:

(i) Decays of cosmic axions which occur after the beginning of the BBN epoch (at around $t \sim 1$ s), but before last scattering (at around $t \sim 10^{13}$ s). These could disrupt nucleosynthesis and affect the abundances of light elements [53].

(ii) Photoproduction (either primary or secondary) from any axion decays that occur between the epoch of electron-positron annihilation (at around $t \sim 10^3$ s) and last scattering. These can lead to observable distortions of the CMB [54].

(iii) Photoproduction (either primary or secondary) from any axion decays that occur after last scattering. These can lead to peaks and other indicative features in the diffuse X-ray and gamma-ray spectra [55], but such features have not been observed by FERMI [56], EGRET [57], COMPTEL [58,59], or any other X-ray or gamma-ray telescope [60–62].

(iv) Entropy production from late axion decays. This can have observational effects on cosmological parameters, such as the rate of cosmic expansion.

(v) Limits from microwave-cavity-detector experiments such as CARRACK [63] and ADMX [64], which search for cosmic axions via their conversion to photons in the presence of strong magnetic fields.

It should be reiterated that the vast majority of the constraints enumerated above are highly model-dependent. The standard energy-dissipation limit from SN1987A [48], for example, provides one of the most stringent limits on the parameter space of a QCD axion. However, these limits are predicated on the assumption that the axion couples to hadrons with significant strength, and that processes such as $NN \rightarrow NNa$ consequently dominate the axion-production rate. A purely photonic axion, on the other hand, lacks such couplings, and hence can only be generated via interactions such as the Primakoff process $e^- \gamma \rightarrow e^- a$, for which the rate is much smaller. As a result, the bounds on $\hat{f}_X$, M_c, and Λ_G for such an axion are considerably weaker than those for a hadronic axion (see, for example, Refs. [52,65] for an analysis of this constraint for a four-dimensional photonic axion). A variety of other constraints, including bounds from monojet searches at hadron colliders and from the requirement that misalignment production of axions dominates over thermal production, also differ markedly depending on whether or not the axion in question couples to hadrons. Still other bounds, such as that from energy loss in white dwarfs [50], depend sensitively on whether or not a given axion couples to leptons.

In Figs. 9 and 10, we display a series of exclusion plots in $(\hat{f}_X, M_c)$ space, taken from Ref. [32], which indicate the regions of parameter space excluded by the considerations enumerated above. The three panels in Fig. 9 correspond to $\Lambda_G = \{1 \text{ GeV}, 1 \text{ TeV}, 100 \text{ TeV}\}$ for the case of a photonic axion with $c_\gamma = 1$, while the three panels in Fig. 10 correspond to the same choices of Λ_G, but for a hadronic axion with $c_\gamma = c_g = 1$. In each case, we have taken $\xi = g_G = \theta = 1$, with $T_{RH} = 5$ MeV and $H_I = 10^{-3}$ GeV; for the hadronic case, we have also assumed that $C_{a\pi}$, $C_{a\pi N}$, etc., take the values given in Eqs. (2.7) and (2.8). The shaded regions in each panel are excluded by the battery of constraints discussed above. The red region is excluded by CAST data, the magenta region by limits on collider processes in which axions appear as missing energy, the purple region by limits on modifications of Newton's law from Eötvös-type experiments, the orange region by limits on distinguishable features in the diffuse extragalactic X-ray and gamma-ray background spectra, the yellow region by observations of the lifetimes of globular-cluster stars, the cyan region by energy-loss limits from supernova SN1987A, the gray region by the model-consistency requirement that $\Lambda_G < f_X$, and the brown region by the upper bound on the dark-matter relic abundance from WMAP. A black, dashed line corresponding to the condition $y = \pi$ has also been included in each panel for reference. Note that each of the exclusion regions shown, with the exception of that from WMAP, differs from the corresponding exclusion region for a four-dimensional axion. The exclusion regions shown are those appropriate for the five-dimensional axion on which our model is based, and are derived in Ref. [32].

The constraints enumerated above for which no exclusion contour has been included in these figures are generally subleading. For example, the applicable constraints from exotic hadron decays [52] are generally far weaker than the constraints from SN1987A, thermal production, etc., for any given choice of parameters. The constraints arising from observational limits on distortions of the CMB are not particularly stringent either, and turn out not to constrain any portion of the model parameter space shown in any of the panels appearing in Figs. 9 and 10. This is because the regions of parameter space in which the Γ_λ are sizeable are those in which $\hat{f}_X$ is quite small, meaning that the Ω_λ are also quite small, as is evident from Fig. 7. Constraints related to the effects of late-decaying a_λ on BBN were not explicitly calculated in Ref. [32]. However, exclusion contours derived from BBN constraints on late-decaying particles are expected to be roughly similar to those derived from CMB constraints, and consequently such constraints are not expected to rule out any additional region of model parameter space not already excluded by other considerations. Limits on the effective equation-of-state parameter $w_{eff}(t)$ are not particularly constraining either. This should come as no surprise, given that we
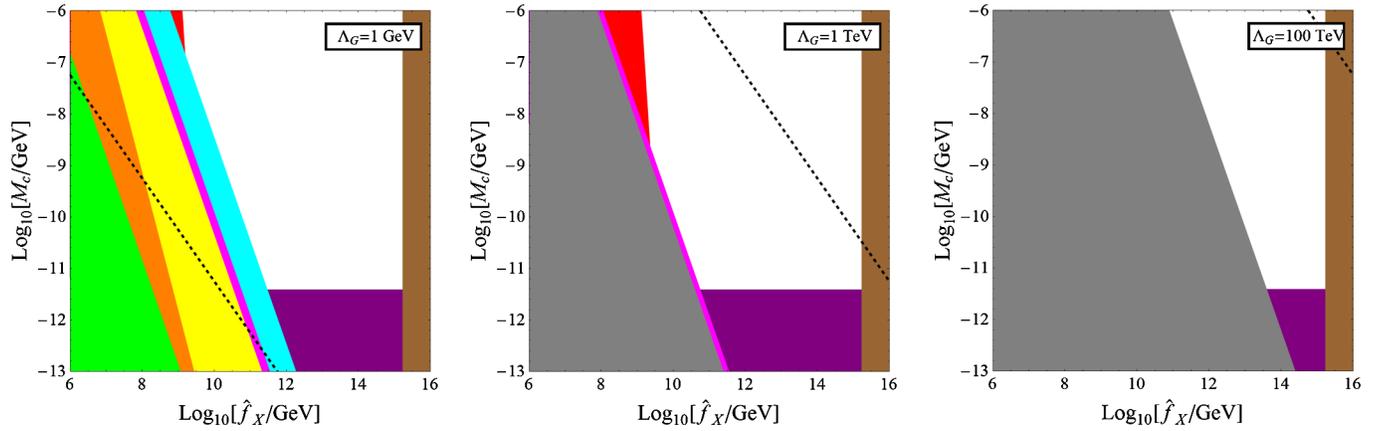
FIG. 9 (color online). Exclusion contours associated with all applicable phenomenological constraints for our bulk-axion model with $\Lambda_G = 1$ GeV (left panel), $\Lambda_G = 1$ TeV (middle panel), and $\Lambda_G = 100$ TeV (right panel). In each case, we have taken $\xi = g_G = 1$, with $T_{RH} = 5$ MeV and $H_I = 10^{-3}$ GeV, and we have assumed that the axion only couples to the photon field. The shaded regions are, respectively, excluded by data from helioscope measurements with CAST (red), collider considerations (magenta), tests of Newton's-law modifications at Eötvös-type experiments (purple), measurements of the diffuse extragalactic X-ray and gamma-ray spectra (orange), observations of the lifetimes of globular-cluster stars (yellow), energy-loss limits from supernova SN1987A (cyan), the model-consistency requirement that $\Lambda_G < f_X$ (gray), and the upper bound on the dark-matter relic abundance from WMAP (brown). The black, dashed line corresponds to the condition $y = \pi$.

showed in Sec. V that the effect of decays on Ω_{tot}^* was negligible within the region of parameter space relevant for dynamical dark matter. A number of additional constraints not listed above also serve to constrain very light axions and axionlike particles [66]; however the particles for which these constraints apply typically involve values of m_X far smaller than those of interest here.

It is evident from these figures that the most stringent constraints on both photonic and hadronic axions are those from SN1987A (cyan) and collider limits on missing-energy processes (magenta). Nevertheless, it is also evident that a hadronic axion is significantly more constrained than a purely photonic axion. As discussed above, the primary reason for this is that the rate of axion production in a thermal setting via interactions with nuclei, pions,

etc., is far larger than the corresponding rate of production via the electron Primakoff process and other interactions which involve the coupling of an axion to photons alone.

Having assessed the phenomenological constraints on a bulk axion, we are now able to definitively address the question as to whether or not our model is a viable model of dynamical dark-matter. In order for this to be so, we require that at least some part of the preferred region in Eq. (5.11) be consistent with the constraints discussed above. Inspecting Figs. 9 and 10, we see that indeed our preferred region is compatible with all of these constraints in both the photonic and hadronic axion cases for $\Lambda_G \gtrsim 100$ GeV, with M_c above the lower bound from Newton's-law modification, but small enough so that
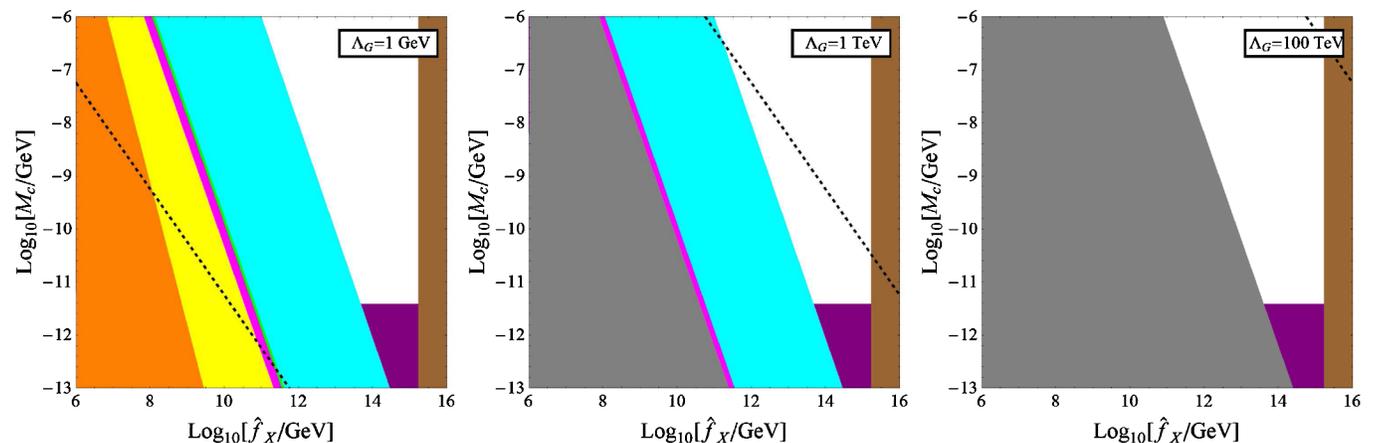

FIG. 10 (color online). Same as in Fig. 9, but for a hadronic axion—i.e., an axion coupled both to the photon and to the gluon field (and hence to pions, nucleons, etc.), but not directly to SM quarks or leptons.

$y \lesssim \pi$. Furthermore, we also see from Figs. 9 and 10 that the phenomenological constraints even permit us to reach deeply into the $y \ll 1$ region. Note that this represents a radical departure from the QCD-axion results presented in Ref. [26]—a departure which is enabled because Λ_G is a free parameter in our model. *Thus, we conclude that within this region of parameter space, our bulk-axion model constitutes a viable, explicit model of dynamical dark matter.*

It should be stated that in addition to the limits discussed above, certain additional astrophysical bounds may also serve to constrain the parameter space of bulk-axion scenarios. For example, it has recently been shown [67] that limits on gamma-ray signals from decaying axions with masses of $\mathcal{O}(10–100 \text{ MeV})$ produced in supernovae can yield an even more stringent limit than that arising from energy-dissipation considerations alone. While these bounds are once again model-dependent (and directly applicable only to cases in which the axion in question couples directly to hadrons with significant strength, and not to a photonic axion), they could provide an important additional constraint on the parameter space of dynamical dark-matter scenarios involving bulk axions. Furthermore, it is also possible that comparable bounds could be obtained from an analysis of photo-emission limits and cooling-rate constraints from neutron stars, similar to that performed for KK gravitons in Ref. [43].

It is important to note that while additional bounds related to axion production in supernovae may serve to further constrain the parameter space of bulk-axion scenarios, these constraints cannot rule out axion models of dynamical dark matter entirely. This is due to the fact that for any given choice of model parameters $\hat{f}_X$, M_c, and Λ_G, the couplings of any mode for which $\lambda \lesssim \pi m_X^2/M_c$ to the SM fields will be suppressed by mixing effects, as discussed in Ref. [13]. Indeed, because we can reach deeply into the $y \ll 1$ region, the magnitude of this coupling suppression can be quite significant. For example, for $y \sim \mathcal{O}(10^{-3})$, we find that the first 20 axion mass eigenmodes have coupling suppressions $\tilde{\lambda}^2 A_\lambda \sim 10^{-6}$. If the coupling suppressions are significant for those a_λ with masses in the "dangerous" range $10 \text{ MeV} \lesssim \lambda \lesssim 100 \text{ MeV}$ discussed above, such a_λ will be produced in supernovae at a negligible rate, and thus all supernova bounds on axion production can be evaded. This can be arranged by demanding that $\pi m_X^2/M_c \gtrsim 1 \text{ GeV}$, so that all modes with masses $\lambda \ll 1 \text{ GeV}$ are effectively in the small-λ regime. Therefore, since Ω_{tot}^* is essentially independent of Λ_G within our preferred region of parameter space, satisfying this condition is simply a matter of choosing a sufficiently large value for Λ_G. Indeed, setting $\hat{f}_X = 10^{14} \text{ GeV}$ in accord with Eq. (5.11), we find that all axion-production constraints from supernovae can be avoided for

$$\Lambda_G \gtrsim (56 \text{ TeV}) \times \left(\frac{M_c}{10^{-11} \text{ GeV}} \right)^{1/4}. \qquad (6.4)$$

We emphasize that this rough bound is not a necessary condition for consistency with supernova data, but a sufficient one. Furthermore, since the neutron-star cooling and photo-emission bounds on KK gravitons rest on the assumption that a population of gravitationally-bound particles of this sort was generated by the supernova whose core-collapse produced a given neutron star, any similar bound on axions would also cease to apply in this regime. We also note that since the fundamental scale f_X is still roughly an order of magnitude larger than the value of Λ_G required to satisfy this condition, given the input values of $\hat{f}_X$ and M_c, no theoretical inconsistency results from positing a confinement scale of this order.

The fact that this coupling-suppression phenomenon is capable of rendering our model consistent with supernova bounds despite the large multiplicity of light modes attests to the importance of this effect in brane/bulk theories. A more detailed overview of this phenomenon and its physical implications will be provided in Ref. [32].

VII. DISCUSSION AND CONCLUSIONS

The aim of this paper has been to present an explicit realization of the dynamical dark-matter framework presented in Ref. [13]. To that end, we have shown that an ensemble consisting of the KK excitations of a light, axion-like field can indeed provide such a realization. Indeed, we have shown that despite the fact that the masses, decay widths, and relic abundances of all of these particles are controlled by only three dimensionful parameters, the ensemble to which they give rise is simultaneously able to reproduce the observed value of Ω_{CDM} and satisfy all applicable constraints from laboratory experiments, astrophysics, and cosmology. As such, this model provides a "proof of concept" for dynamical dark matter as a viable alternative framework for dark-matter physics. In addition, it also provides a method of addressing the dark-matter question which does not require the introduction of any additional stabilizing symmetry.

Many qualifications, extensions, and possible generalizations of our dynamical dark-matter framework were discussed at the end of Ref. [13]; here, we shall restrict our attention to five points which are specific to the bulk-axion model presented in this paper.

(i) First, in this work, we have made use of the rapid-turn-on approximation in Eq. (4.8) in calculating the relic abundances of the a_λ. As discussed in Sec. IV, this approximation is well motivated, since the instanton-generated mass term $m_X(T)$ falls rapidly with temperature when $T \gtrsim \Lambda_G$. Furthermore, the primary results of this paper are essentially insensitive to this approximation. This is because the fields which contribute significantly to Ω_{CDM} in regions of parameter space which yield a realistic dark-matter

relic abundance begin oscillating only well after $m_X(T)$ has already settled into its constant, late-time value. However, the relic abundance of any field which begins oscillating before $m_X(T)$ takes this late-time value will, in general, depend on the details of how this mass evolves in time. The quantitative effect on the abundance of a single field has long been appreciated [22], but in our model, the effects are more complicated and more subtle because we have a coupled system of mixed scalars with different masses and therefore different oscillation times. It would be interesting to examine how a more rigorous treatment of the turn-on of $m_X(T)$ would affect Ω_{tot} and η in situations in which these quantities are sensitive to the time-dependence of this brane-mass term. Such a study would have important implications for more general scenarios involving other kinds of light bulk scalars. Indeed, the relationship between the size of the brane-mass term for such scalars and the time at which that brane mass is dynamically generated may differ significantly from the relationship which holds for axions.

(ii) Second, as alluded to in Sec. VI, it may be possible to further test or constrain the parameter space of bulk-axion models of dynamical dark matter in a number of ways. We have already mentioned one potential constraint which derives from limits on high-energy photons resulting from the decays of axions produced in supernovae [67]. Other considerations also merit investigation. For example, a detailed analysis of the limits imposed by BBN on scenarios involving multiple decaying fields with different lifetimes and abundances could provide important constraints on dynamical dark-matter models in general. In addition, other considerations, such as limits on mass loss and decreases in the dark-matter density in the halos of dwarf galaxies [68], could also be used to constrain dynamical dark-matter models. Indeed, while a number of standard constraints on individual unstable relic particles in the early universe have been revisited in a dynamical dark-matter context [32], it would be interesting to see how other constraints would apply in this context as well.

(iii) Third, we note that we have not specified a particular model of inflation as part of the cosmological context for our model. Indeed, other than requiring a low reheating temperature $T_{\text{RH}} \sim \mathcal{O}(\text{MeV})$, we have remained largely agnostic about the details of the inflationary model, the form of the inflaton potential, or even the scale H_I. For the most part, our model does not depend on these particulars. However, certain consistency conditions do place meaningful restrictions on the set of inflationary scenarios with which our model is compatible. One such condition can be derived from the fact that vacuum fluctuations during inflation generically give rise to a background value $\langle \phi^2 \rangle \approx H_I^3 t_I / 4\pi^2$ for any scalar ϕ with a mass $m_\phi \ll H_I$, where t_I is the duration of inflation. This implies that the relationship between the mass λ and initial energy density ρ_λ in Eq. (4.12) in our model is truly valid only for the lighter a_λ in a given tower—*i.e.*, those for which $\theta^2 A_\lambda^2 \hat{f}_X^2 \gtrsim H_I^3 t_I / 4\pi^2$. By contrast, any heavier a_λ which still satisfy $\lambda \ll H_I$ receive the leading contributions to their background values from vacuum fluctuations during inflation, and thus effectively acquire an initial abundance $\rho_\lambda \sim \lambda^2 H_I^3 t_I$. In typical scenarios, we expect $H_I t_I \approx N_e \sim \mathcal{O}(60)$, where N_e is the number of e-foldings of inflation. The results for Ω_{tot} derived in Sec. V therefore remain consistent, provided that $\hat{f}_X^2 \gg H_I^2$. Indeed, since $\hat{f}_X \sim 10^{14}$–10^{15} GeV within the preferred region of parameter space specified in Eq. (5.11), we see that $\hat{f}_X \gg H_I$ is certainly not inconsistent with our model and is in fact even expected. However, this condition on H_I has non-trivial implications for inflationary models. While a low scale for H_I is certainly not excluded (see, e.g., Refs. [39,69]), extremely small values of H_I tend to be rather nongeneric [70] among typical classes of inflationary potentials, and thus require either substantial tuning or careful construction. Indeed, any consistent inflationary model of this sort must give rise to density fluctuations on a scale consistent with constraints from CMB data [2], such as those on the spectral index n_s, and must also satisfy other observational constraints. The development of explicit inflationary scenarios of this sort is therefore an interesting topic for future investigation.

(iv) Fourth, we note that while we have chosen in this paper to focus on the case in which the ensemble of fields reproducing Ω_{CDM} are the KK excitations of a bulk-axion field, such a field is by no means unique in possessing the characteristics necessary to give rise to such an ensemble. Indeed, as discussed in Ref. [13], much of the analysis presented here pertains to any light bulk scalar for which a mass term is dynamically generated via its interactions with brane-localized fields. Furthermore, for a generic bulk scalar, the relationship between the time at which this mass term is dynamically generated and the magnitude of this mass term itself may differ from that which relates t_G and m_X for a bulk axion. As a result, much more freedom may exist for constructing viable models within the dynamical dark-matter framework. For example, light moduli could also, in principle, provide a viable model of dynamical dark matter.

(v) Finally, we emphasize that the presence of additional axionlike fields is fairly generic, and perhaps

even expected, in many theoretically motivated scenarios for physics beyond the standard model (see, e.g., Ref. [25]). Moreover, it has even been argued that many of these axionlike fields are likely to be light [66]. Thus, the discovery of a vast ensemble of axionlike particles could provide important insight into what physics looks like at high scales. Indeed, if many of these axions have relatively small masses, we find ourselves in the intriguing situation in which most of the matter in the universe is simultaneously both light and dark.

Our goal in this work has been to provide an existence proof for dynamical dark matter—i.e., to provide a model in which lifetimes are balanced against abundances in such a way that the ensemble of dark-matter particles successfully reproduces Ω_{CDM} while at the same time satisfying all phenomenological constraints. As we have seen in this paper, our bulk-axion model indeed passes this test. In one sense, our model does so in the most interesting way possible: with $y \ll 1$ (signifying that our tower of axion KK modes is highly mixed) and with a tower fraction η which is significantly different from zero. In another sense, however, this model is fairly conservative: those modes which contribute most to Ω_{tot}^* turn out to be rather long-lived, and likewise our numerical result for w_* within the preferred region of parameter space turns out to be rather close to zero. Indeed, at first glance, one might suspect that these latter properties are in fact generic for dynamical dark-matter models, or even that such models are therefore really no different from traditional dark-matter models in terms of their abundance and stability requirements.

This is not the case, however, for the balancing of lifetimes against abundances—which is the hallmark of the dynamical dark-matter framework—is precisely why this framework does *not* require such a degree of stability, much less the existence of a stabilizing symmetry. While certain accidental features of our bulk-axion model result in a preferred region of parameter space which is somewhat conservative, we emphasize that these features are not generic even to theories with bulk scalars, much less realistic dynamical dark-matter models as a whole. Note, for example, that a particular relationship exists in bulk-axion models between the mass λ of a given KK mass eigenstate a_λ, the strength of its effective coupling to SM fields, and the overall magnitude of its relic abundance Ω_λ through the dependence of these quantities on $\hat{f}_X$. Even for other bulk scalars (*e.g.*, moduli), these relationships do not necessarily hold. There is therefore no reason to expect dynamical dark-matter models based around such fields to be as conservative as the axion model we have presented here.

In this connection, there is an even more important point that deserves emphasis. In dynamical dark-matter scenarios, we have no single characteristic decay width Γ nor abundance Ω, but rather an entire *spectrum* of widths Γ_λ and abundances Ω_λ. This therefore begs the fundamental question: if our "proof of concept" model presented here is to be viewed as somewhat conservative, how far from the conservative limit can we go?

At first glance, one might try to answer this question by attempting to determine, for each time t during the evolution of the universe, the maximum abundance $\Omega_{max}(t)$ that a given component in a dark-matter ensemble may have if it has a lifetime $\tau \sim t$. In other words, given the entirety of the cosmological constraints from BBN, CMB distortions, etc., there exists a *function* $\Omega_{max}(\Gamma)$ which describes the maximum abundance any dark-matter constituent may have as a function of its decay width. It might therefore seem that knowledge of this function would uniquely determine the full range of possibilities inherent in our dynamical dark-matter framework.

Such an approach to answering our fundamental question is, in a sense, already a departure from the usual manner of approaching dark-matter physics. However, even the notion of such a function $\Omega_{max}(\Gamma)$ relies too strongly on a single-particle perspective. One of the critical features of our dynamical dark-matter framework is that it involves a vast *ensemble* of dark-matter components. Some of these components might decay earlier in cosmological evolution, while others might decay later. As a result, the maximum abundance that a given component may have if it decays on a characteristic time scale τ will itself be directly affected not only by the abundances of all of the other components with earlier characteristic decay times $\tau' < \tau$, but even the components with $\tau' > \tau$. Moreover, as we have seen, most phenomenological constraints on dark-matter decays are sensitive not merely to what happens at a specific moment in time, but to the integrated effects of such decays over a broad range of time scales. In other words, our dynamical dark-matter framework teaches us that astrophysical and cosmological constraints do not lead to a single function $\Omega_{max}(\Gamma)$, but rather a more subtle set of intertwined constraints on lifetimes and abundances across our entire dark-matter ensemble as a whole.

Clearly, this issue has not been studied in any detail in the literature. However, it is readily apparent that this is indeed the only proper way in which one should express constraints on particle decays from a generic dark sector. Viewed from this perspective, then, the existence of even one viable dynamical dark-matter model—no matter how "conservative" it might be—gives us strong motivation to reexamine cosmological and astrophysical constraints within this framework. Indeed, it is only in this way that we will be able to fully explore our dynamical dark-matter framework, and understand its full range of phenomenological possibilities.

APPENDIX. EVOLUTION OF A DECAYING AXION FIELD

For completeness, in this Appendix we provide exact solutions to Eq. (4.6) for a real-valued function $a_\lambda(t)$. These take the form

$$a_\lambda(t) = c_\lambda^{(M)} \tilde{M}_\kappa(\lambda, t) + c_\lambda^{(U)} \tilde{U}_\kappa(\lambda, t), \qquad (A1)$$

where $c_\lambda^{(M)}$ and $c_\lambda^{(U)}$ are undetermined constants, and

$$\tilde{M}_\kappa(\lambda, t) \equiv e^{-(k_\lambda + \Gamma_\lambda)t/2}\left[M\left(\frac{\kappa(k_\lambda + \Gamma_\lambda)}{2k_\lambda}, \kappa, k_\lambda t\right) + e^{k_\lambda t} M\left(\frac{\kappa(k_\lambda - \Gamma_\lambda)}{2k_\lambda}, \kappa, -k_\lambda t\right) \right]$$

$$\tilde{U}_\kappa(\lambda, t) \equiv e^{-(k_\lambda + \Gamma_\lambda)t/2}\left[U\left(\frac{\kappa(k_\lambda + \Gamma_\lambda)}{2k_\lambda}, \kappa, k_\lambda t\right) + e^{k_\lambda t} U\left(\frac{\kappa(k_\lambda - \Gamma_\lambda)}{2k}, \kappa, -k_\lambda t\right) \right]. \qquad (A2)$$

In these expressions, $k_\lambda \equiv \sqrt{\Gamma_\lambda^2 - 4\lambda^2}$, κ was defined in Eq. (4.7), $M(a, b, x)$ denotes Kummer's confluent hypergeometric function

$$M(a, b, x) = \sum_{m=0}^{\infty} \frac{(a)_m x^m}{(b)_m m!}, \qquad (A3)$$

where $(x)_n = (x + n - 1)!/(x - 1)!$ is the Pochhammer function, and $U(a, b, x)$ denotes the Tricomi confluent hypergeometric function

$$U(a, b, x) = \frac{\Gamma(1 - b)}{\Gamma(a - b + 1)} M(a, b, x) + \frac{\Gamma(b - 1)}{\Gamma(a)} x^{1-b} M(a - b + 1, 2 - b, x). \qquad (A4)$$

It can be verified upon setting $\Gamma_\lambda = 0$ that Eq. (A1) reduces to the exact form obtained for a tower of stable KK axions in Ref. [26].

The values of $c_\lambda^{(M)}$ and $c_\lambda^{(U)}$ in Eq. (A1) are determined by the initial conditions chosen for $a_\lambda(t)$ and $\dot{a}_\lambda(t)$ at $t = t_0$, where t_0 is some initial time. Expressed in terms of these initial values, this equation takes the general form

$$a_\lambda(t) = \frac{[\dot{a}_\lambda(t_0)\tilde{U}_\kappa(\lambda, t_0) - a_\lambda(t_0)\dot{\tilde{U}}_\kappa(\lambda, t_0)]\tilde{M}_\kappa(\lambda, t) - [\dot{a}_\lambda(t_0)\tilde{M}_\kappa(\lambda, t_0) - a_\lambda(t_0)\dot{\tilde{M}}_\kappa(\lambda, t_0)]\tilde{U}_\kappa(\lambda, t)}{\dot{\tilde{M}}_\kappa(\lambda, t_0)\tilde{U}_\kappa(\lambda, t_0) - \dot{\tilde{U}}_\kappa(\lambda, t_0)\tilde{M}_\kappa(\lambda, t_0)}, \qquad (A5)$$

where the time derivatives of $\tilde{M}_\kappa(\lambda, t)$ and $\tilde{U}_\kappa(\lambda, t)$ have the explicit forms

$$\dot{\tilde{M}}_\kappa(\lambda, t) = e^{-(k_\lambda + \Gamma_\lambda)t/2}(k_\lambda - \Gamma_\lambda)\left[M\left(\frac{\kappa(k_\lambda + \Gamma_\lambda)}{2k_\lambda}, \kappa, k_\lambda t\right) - e^{k_\lambda t} M\left(\frac{\kappa(k_\lambda - \Gamma_\lambda)}{2k_\lambda} + 1, \kappa + 1, -k_\lambda t\right) \right]$$

$$\dot{\tilde{U}}_\kappa(\lambda, t) = -\frac{1}{2} e^{-(k_\lambda + \Gamma_\lambda)t/2}\Bigg\{ (k_\lambda + \Gamma_\lambda)\left[U\left(\frac{\kappa(k_\lambda + \Gamma_\lambda)}{2k_\lambda}, \kappa, k_\lambda t\right) + \kappa U\left(\frac{\kappa(k_\lambda + \Gamma_\lambda)}{2k_\lambda} + 1, \kappa + 1, k_\lambda t\right) \right]$$

$$- e^{k_\lambda t}(k_\lambda - \Gamma_\lambda)\left[U\left(\frac{\kappa(k_\lambda - \Gamma_\lambda)}{2k_\lambda}, \kappa, -k_\lambda t\right) + \kappa U\left(\frac{\kappa(k_\lambda - \Gamma_\lambda)}{2k_\lambda} + 1, \kappa + 1, -k_\lambda t\right) \right] \Bigg\}. \qquad (A6)$$

Once again, if we set $\Gamma_\lambda = 0$ in this expression (which also implies that $k_\lambda = 2i\lambda$), we recover the result

$$a_\lambda(t) \xrightarrow{\Gamma_\lambda \to 0} -\frac{\pi}{\sqrt{2}} a_\lambda(t_0)\lambda t_0^{5/4} t^{-1/4}[J_{-5/4}(\lambda t_0)J_{1/4}(\lambda t) + J_{5/4}(\lambda t_0)J_{-1/4}(\lambda t)], \qquad (A7)$$

which agrees with the result obtained in Ref. [26].

In the rapid-turn-on approximation, in which $m_X(t)$ takes the Heaviside form specified in Eq. (4.8), the initial conditions for a_λ and $\dot{a}_\lambda$ at $t_0 = t_\lambda$ take the form given in Eq. (4.9). Upon substituting these initial conditions into Eq. (A5), we find that during the cosmological epoch in which coherent oscillations of a given a_λ begin, we have

$$a_\lambda(t) = \theta \hat{f}_X A_\lambda \frac{\dot{\tilde{M}}_\kappa(\lambda, t_\lambda)\tilde{U}_\kappa(\lambda, t) - \dot{\tilde{U}}_\kappa(\lambda, t_\lambda)\tilde{M}_\kappa(\lambda, t)}{\dot{\tilde{M}}_\kappa(\lambda, t_\lambda)\tilde{U}_\kappa(\lambda, t_\lambda) - \dot{\tilde{U}}_\kappa(\lambda, t_\lambda)\tilde{M}_\kappa(\lambda, t_\lambda)}. \qquad (A8)$$

The value of any a_λ during subsequent epochs can then be obtained iteratively from this relation.

[1] G. Jungman, M. Kamionkowski, and K. Griest, Phys. Rep. **267**, 195 (1996); K. A. Olive, arXiv:astro-ph/0301505; D. Hooper, arXiv:0901.4090; N. Weiner, "Dark Matter Theory," video lectures given at TASI 2009, http://physicslearning2.colorado.edu/tasi/tasi_2009/tasi_2009.htm; J. L. Feng, Annu. Rev. Astron. Astrophys. **48**, 495 (2010).

[2] E. Komatsu *et al.* (WMAP Collaboration), Astrophys. J. Suppl. Ser. **180**, 330 (2009).

[3] X. L. Chen and M. Kamionkowski, Phys. Rev. D **70**, 043502 (2004); L. Zhang, X. Chen, M. Kamionkowski, Z. G. Si, and Z. Zheng, Phys. Rev. D **76**, 061301 (2007).

[4] C. Boehm, P. Fayet, and J. Silk, Phys. Rev. D **69**, 101302 (2004); E. Ma, Annales Fond. Broglie **31**, 285 (2006); T. Hur, H. S. Lee, and S. Nasri, Phys. Rev. D **77**, 015008 (2008); M. Adibzadeh and P. Q. Hung, Nucl. Phys. **B804**, 223 (2008); J. L. Feng and J. Kumar, Phys. Rev. Lett. **101**, 231301 (2008); H. Sung Cheon, S. K. Kang, and C. S. Kim, Phys. Lett. B **675**, 203 (2009); J. H. Huh, J. E. Kim, and B. Kyae, Phys. Rev. D **79**, 063529 (2009); M. Fairbairn and J. Zupan, J. Cosmol. Astropart. Phys. 07 (2009) 001; K. M. Zurek, Phys. Rev. D **79**, 115002 (2009); H. Baer, M. Haider, S. Kraml, S. Sekmen, and H. Summy, J. Cosmol. Astropart. Phys. 02 (2009) 002; B. Batell, M. Pospelov, and A. Ritz, Phys. Rev. D **79**, 115019 (2009); S. Profumo, K. Sigurdson, and L. Ubaldi, J. Cosmol. Astropart. Phys. 12 (2009) 016; F. Chen, J. M. Cline, and A. R. Frey, Phys. Rev. D **80**, 083516 (2009); H. Zhang, C. S. Li, Q. H. Cao, and Z. Li, Phys. Rev. D **82**, 075003 (2010); I. Cholis and N. Weiner, arXiv:0911.4954; X. Gao, Z. Kang, and T. Li, Eur. Phys. J. C **69**, 467 (2010); F. D'Eramo and J. Thaler, J. High Energy Phys. 06 (2010) 109; D. Feldman, Z. Liu, P. Nath, and G. Peim, Phys. Rev. D **81**, 095017 (2010).

[5] P. T. Winslow, K. Sigurdson, and J. N. Ng, Phys. Rev. D **82**, 023512 (2010).

[6] G. Servant and T. M. P. Tait, Nucl. Phys. **B650**, 391 (2003); H. C. Cheng, J. L. Feng, and K. T. Matchev, Phys. Rev. Lett. **89**, 211301 (2002).

[7] I. Antoniadis, Phys. Lett. B **246**, 377 (1990); I. Antoniadis, K. Benakli, and M. Quiros, Phys. Lett. B **331**, 313 (1994).

[8] K. R. Dienes, E. Dudas, and T. Gherghetta, Phys. Lett. B **436**, 55 (1998); Nucl. Phys. **B537**, 47 (1999); arXiv:hep-ph/9807522.

[9] T. Appelquist, H. C. Cheng, and B. A. Dobrescu, Phys. Rev. D **64**, 035002 (2001).

[10] H. C. Cheng and I. Low, J. High Energy Phys. 09 (2003) 051; 08 (2004) 061.

[11] N. Arkani-Hamed, A. G. Cohen, and H. Georgi, Phys. Lett. B **513**, 232 (2001).

[12] N. Cabibbo, G. R. Farrar, and L. Maiani, Phys. Lett. B **105**, 155 (1981); P. Salati and J. C. Wallet, Phys. Lett. B **144**, 61 (1984); M. S. Turner, G. Steigman, and L. M. Krauss, Phys. Rev. Lett. **52**, 2090 (1984); G. Gelmini, D. N. Schramm, and J. W. F. Valle, Phys. Lett. B **146**, 311 (1984); A. G. Doroshkevich and M. Y. Khlopov, Yad. Fiz. **39**, 869 (1984); Mon. Not. R. Astron. Soc. **211**, 279 (1984); Pis'ma Astron. Zh. **11**, 563 (1985); A. G. Doroshkevich, A. A. Klypin, and M. Y. Khlopov, Astron. Zh. **65**, 248 (1988); A. G. Doroshkevich, M. Khlopov, and A. A. Klypin, Mon. Not. R. Astron. Soc. **239**, 923 (1989);

J. R. Ellis, J. L. Lopez, and D. V. Nanopoulos, Phys. Lett. B **247**, 257 (1990); Z. G. Berezhiani and M. Y. Khlopov, Z. Phys. C **49**, 73 (1991); Yad. Fiz. **52**, 96 (1990) [Sov. J. Nucl. Phys. **52**, 60 (1990)]; V. Berezinsky, A. Masiero, and J. W. F. Valle, Phys. Lett. B **266**, 382 (1991); S. Dodelson and J. M. Jubas, Mon. Not. R. Astron. Soc. **266**, 886 (1994); T. Asaka, J. Hashiba, M. Kawasaki, and T. Yanagida, Phys. Rev. D **58**, 023507 (1998); D. J. H. Chung, E. W. Kolb, and A. Riotto, Phys. Rev. D **59**, 023501 (1998); A. G. Doroshkevich and P. D. Naselsky, Phys. Rev. D **65**, 123517 (2002); H. B. Kim and J. E. Kim, Phys. Lett. B **527**, 18 (2002); S. H. Hansen and Z. Haiman, Astrophys. J. **600**, 26 (2004); J. L. Feng, A. Rajaraman, and F. Takayama, Phys. Rev. Lett. **91**, 011302 (2003); Phys. Rev. D **68**, 063504 (2003); X. J. Bi, M. Z. Li, and X. M. Zhang, Phys. Rev. D **69**, 123521 (2004).

[13] K. R. Dienes and B. Thomas, Phys. Rev. D **85**, 083523 (2012).

[14] E. W. Kolb and M. S. Turner, Front. Phys. **69**, 1 (1990).

[15] R. D. Peccei, Lect. Notes Phys. **741**, 3 (2008).

[16] J. E. Kim and G. Carosi, Rev. Mod. Phys. **82**, 557 (2010).

[17] M. P. Hertzberg, M. Tegmark, and F. Wilczek, Phys. Rev. D **78**, 083507 (2008).

[18] R. D. Peccei and H. R. Quinn, Phys. Rev. Lett. **38**, 1440 (1977); Phys. Rev. D **16**, 1791 (1977).

[19] C. A. Baker *et al.*, Phys. Rev. Lett. **97**, 131801 (2006).

[20] S. Weinberg, Phys. Rev. Lett. **40**, 223 (1978); F. Wilczek, Phys. Rev. Lett. **40**, 279 (1978).

[21] D. J. Gross, R. D. Pisarski, and L. G. Yaffe, Rev. Mod. Phys. **53**, 43 (1981).

[22] M. S. Turner, Phys. Rev. D **33**, 889 (1986).

[23] S. Chang and K. Choi, Phys. Lett. B **316**, 51 (1993).

[24] J. E. Kim, Phys. Rev. Lett. **43**, 103 (1979); M. A. Shifman, A. I. Vainshtein, and V. I. Zakharov, Nucl. Phys. B **166**, 493 (1980).

[25] E. Witten, Phys. Lett. B **155**, 151 (1985).

[26] K. R. Dienes, E. Dudas, and T. Gherghetta, Phys. Rev. D **62**, 105023 (2000).

[27] N. Arkani-Hamed, S. Dimopoulos, and G. R. Dvali, Phys. Lett. B **429**, 263 (1998).

[28] N. Arkani-Hamed, S. Dimopoulos, and G. R. Dvali, Phys. Rev. D **59**, 086004 (1999).

[29] L. Randall and R. Sundrum, Phys. Rev. Lett. **83**, 3370 (1999); **83**, 4690 (1999).

[30] E. Witten, Nucl. Phys. **B471**, 135 (1996); J. D. Lykken, Phys. Rev. D **54**, R3693 (1996); I. Antoniadis, N. Arkani-Hamed, S. Dimopoulos, and G. R. Dvali, Phys. Lett. B **436**, 257 (1998); G. Shiu and S. H. H. Tye, Phys. Rev. D **58**, 106007 (1998); C. P. Bachas, J. High Energy Phys. 11 (1998) 023.

[31] T. Flacke, B. Gripaios, J. March-Russell, and D. Maybury, J. High Energy Phys. 01 (2007) 061; T. Flacke and D. Maybury, J. High Energy Phys. 03 (2007) 007.

[32] K. R. Dienes and B. Thomas, arXiv:1203.1923.

[33] K. Nakamura *et al.* (Particle Data Group), J. Phys. G **37**, 075021 (2010).

[34] D. J. Kapner, T. S. Cook, E. G. Adelberger, J. H. Gundlach, B. R. Heckel, C. D. Hoyle, and H. E. Swanson, Phys. Rev. Lett. **98**, 021101 (2007).

[35] C. Macesanu and M. Trodden, Phys. Rev. D **71**, 024008 (2005).

[36] L. J. Hall and D. Tucker-Smith, Phys. Rev. D **60**, 085008 (1999).

[37] P. Binetruy, C. Deffayet, and D. Langlois, Nucl. Phys. **B565**, 269 (2000).

[38] M. Kawasaki, K. Kohri, and N. Sugiyama, Phys. Rev. Lett. **82**, 4168 (1999); Phys. Rev. D **62**, 023506 (2000); K. Ichikawa, M. Kawasaki, and F. Takahashi, Phys. Rev. D **72**, 043522 (2005).

[39] D. Grin, T. L. Smith, and M. Kamionkowski, Phys. Rev. D **77**, 085020 (2008).

[40] L. Visinelli and P. Gondolo, Phys. Rev. D **81**, 063508 (2010).

[41] S. Chang, S. Tazawa, and M. Yamaguchi, Phys. Rev. D **61**, 084005 (2000).

[42] C. Bambi, M. Kawasaki, and F. R. Urban, Phys. Rev. D **80**, 023533 (2009).

[43] S. Hannestad and G. G. Raffelt, Phys. Rev. Lett. **88**, 071301 (2002).

[44] C. Hanhart, J. A. Pons, D. R. Phillips, and S. Reddy, Phys. Lett. B **509**, 1 (2001).

[45] S. Hannestad and G. Raffelt, Phys. Rev. Lett. **87**, 051301 (2001).

[46] D. M. Lazarus, G. C. Smith, R. Cameron, A. C. Melissinos, G. Ruoso, Y. K. Semertzidis, and F. A. Nezrick, Phys. Rev. Lett. **69**, 2333 (1992).

[47] J. Jaeckel and A. Ringwald, arXiv:1002.0329.

[48] G. G. Raffelt, Lect. Notes Phys. **741**, 51 (2008).

[49] P. Gondolo and G. Raffelt, Phys. Rev. D **79**, 107301 (2009).

[50] G. G. Raffelt, Phys. Lett. B **166**, 402 (1986).

[51] G. F. Giudice, R. Rattazzi, and J. D. Wells, Nucl. Phys. **B544**, 3 (1999); T. Han, J. D. Lykken, and R. J. Zhang, Phys. Rev. D **59**, 105006 (1999); G. F. Giudice, T. Plehn, and A. Strumia, Nucl. Phys. **B706**, 455 (2005).

[52] E. Masso and R. Toldra, Phys. Rev. D **52**, 1755 (1995).

[53] R. H. Cyburt, J. R. Ellis, B. D. Fields, and K. A. Olive, Phys. Rev. D **67**, 103521 (2003); M. Kawasaki, K. Kohri, and T. Moroi, Phys. Lett. B **625**, 7 (2005); Phys. Rev. D **71**, 083502 (2005).

[54] W. Hu and J. Silk, Phys. Rev. Lett. **70**, 2661 (1993); Phys. Rev. D **48**, 485 (1993).

[55] M. Kawasaki and T. Yanagida, Phys. Lett. B **399**, 45 (1997).

[56] A. A. Abdo et al. (The Fermi-LAT collaboration), Phys. Rev. Lett. **104**, 101101 (2010).

[57] A. W. Strong, I. V. Moskalenko, and O. Reimer, Astrophys. J. **613**, 956 (2004).

[58] S. C. Kappadath et al., Bull. Am. Astron. Soc. **30**, 926 (1998); Ph.D. thesis, http://www.gro.sr.unh.edu/users/ckappada/ckappada.html.

[59] P. Sreekumar, F. W. Stecker, and S. C. Kappadath, AIP Conf. Proc. **510**, 459 (2000).

[60] R. L. Kinzer, G. V. Jung, D. E. Gruber, J. L. Matteson, and L. E. Peterson, Astrophys. J. **475**, 361 (1997); D. E. Gruber, J. L. Matteson, L. E. Peterson, and G. V. Jung, arXiv:astro-ph/9903492.

[61] R. C. Hickox and M. Markevitch, Astrophys. J. **645**, 95 (2006).

[62] W. N. Brandt et al., Astron. J. **122**, 2810 (2001); R. Giacconi et al., Astrophys. J. Suppl. Ser. **139**, 369 (2002).

[63] K. Yamamoto et al., arXiv:hep-ph/0101200.

[64] S. J. Asztalos et al. (The ADMX Collaboration), Phys. Rev. Lett. **104**, 041301 (2010).

[65] E. Masso and R. Toldra, Phys. Rev. D **55**, 7967 (1997).

[66] A. Arvanitaki, S. Dimopoulos, S. Dubovsky, N. Kaloper, and J. March-Russell, Phys. Rev. D **81**, 123530 (2010).

[67] M. Giannotti, L. D. Duffy, and R. Nita, J. Cosmol. Astropart. Phys. 01 (2011) 015.

[68] A. H. G. Peter, C. E. Moody, A. J. Benson, and M. Kamionkowski, Proc. Sci. IDM2010 (2011) 084.

[69] L. Randall and S. D. Thomas, Nucl. Phys. **B449**, 229 (1995); G. German, G. G. Ross, and S. Sarkar, Nucl. Phys. **B608**, 423 (2001); M. Giovannini, Phys. Rev. D **67**, 123512 (2003); K. Dimopoulos, D. H. Lyth, and Y. Rodriguez, J. High Energy Phys. 02 (2005) 055; P. Q. Hung, E. Masso, and G. Zsembinszki, J. Cosmol. Astropart. Phys. 12 (2006) 004; R. Allahverdi, K. Enqvist, J. Garcia-Bellido, A. Jokinen, and A. Mazumdar, J. Cosmol. Astropart. Phys. 06 (2007) 019; R. Allahverdi, B. Dutta, and K. Sinha, Phys. Rev. D **81**, 083538 (2010); G. G. Ross and G. German, Phys. Lett. B **691**, 117 (2010).

[70] L. A. Boyle, P. J. Steinhardt, and N. Turok, Phys. Rev. Lett. **96**, 111301 (2006).