

Towards string theory expectations for photon couplings to axionlike particles

James Halverson, Cody Long, Brent Nelson, and Gustavo Salinas 

Department of Physics, Northeastern University Boston, Massachusetts 02115-5000, USA



(Received 25 September 2019; published 21 November 2019)

Axionlike particle (ALP)–photon couplings are modeled in large ensembles of string vacua and random matrix theories. In all cases, the effective coupling increases polynomially in the number of ALPs, of which hundreds or thousands are expected in the string ensembles, many of which are ultralight. The expected value of the couplings $g_{a\gamma\gamma} \simeq 10^{-12} \text{ GeV}^{-1} - 10^{-10} \text{ GeV}^{-1}$ provide viable targets for future x-ray telescopes and axion helioscopes, and in some cases are already in tension with existing data.

DOI: [10.1103/PhysRevD.100.106010](https://doi.org/10.1103/PhysRevD.100.106010)

I. INTRODUCTION

If string theory is the correct theory of quantum gravity, one of its vacua must realize the photon of classical electromagnetism. Uncharged spin zero particles must couple to the electromagnetic field strength, since all couplings in string theory are determined by vacuum expectation values (VEVs) of scalar fields. This applies not only to the usual parity even operator $F_{\mu\nu}F^{\mu\nu}$, but also the parity odd operator, requiring the existence of a coupling

$$\mathcal{L} \supset -\frac{1}{4}g_{a\gamma\gamma}aF_{\mu\nu}\tilde{F}^{\mu\nu} \quad (1)$$

in the effective Lagrangian, where $\tilde{F}^{\mu\nu} = \epsilon^{\mu\nu\rho\sigma}F_{\rho\sigma}$. The pseudoscalar a is an axionlike particle (ALP), which is not necessarily the QCD axion, and it may be a nontrivial linear combination of the hundreds or thousands of ALPs expected from studies of string vacua.

Numerous ground-based experiments [1–4] and satellite observations [5–11] place constraints on ALP-photon interactions, probing widely different regimes for the axion mass m_a and coupling strength $g_{a\gamma\gamma}$. Existing limits are already remarkable, within a few orders of magnitude of grand unified theory (GUT) scale decay constants $f_{a\gamma\gamma} \equiv g_{a\gamma\gamma}^{-1}$. However, the exclusions depend critically on whether the ALP is assumed to be a sizable fraction of the dark matter and also experimental limitations affecting the accessible mass range. Since dark matter in string theory is often multicomponent (e.g., [12,13]) and many ALPs are expected to be ultralight, but not yet in any fixed mass

window, we will focus on analyses relevant for experiments that do not require ALP dark matter and only set limits below a fixed mass threshold.

For such experiments, the strongest bounds on the coupling arise at low mass, $m_a \leq 10^{-2} \text{ eV}$, where results [1] from the axion helioscope CAST require $g_{a\gamma\gamma} \leq 7 \times 10^{-11} \text{ GeV}^{-1}$. The CAST result already significantly outperforms projected collider bounds for $m_a \simeq \text{GeV}$ scale ALPs, leading us to focus on the low mass range. For even lower masses, $m_a \leq 10^{-12} \text{ eV}$, observations from the x-ray telescope Chandra require $g_{a\gamma\gamma} \leq 8 \times 10^{-13} \text{ GeV}^{-1}$ [10]. The future helioscope IAXO and satellite STROBE-X are projected to probe $g_{a\gamma\gamma} \simeq 2 \times 10^{-12} \text{ GeV}^{-1}$ for $m_a \leq 10^{-2} \text{ eV}$ [3] and $g_{a\gamma\gamma} \simeq 8 \times 10^{-14} \text{ GeV}^{-1}$ for $m_a \leq 10^{-12} \text{ eV}$ [14], respectively.

The primary focus of this work is to understand the dependence of $g_{a\gamma\gamma}$ on the number N of ALPs in both models of string vacua and random matrix effective field theories. Though in the $N = 1$ case (using mild assumptions described in the text) we have

$$\text{Single ALP: } g_{a\gamma\gamma} = \frac{1}{\sqrt{3}M_p}, \quad (2)$$

independent of the details of the compactification geometry (including the string scale), it is reasonable to imagine that the introduction of additional ALPs will increase the effective coupling to photons. Concretely, we study whether the bulk of the $g_{a\gamma\gamma}$ distribution at large N is in a range probed by current or future experiments. This is of interest because most vacua are expected to arise at the largest values of N afforded by the given ensemble. We will refer to the mean of the distribution at this value of N as the expected value of $g_{a\gamma\gamma}$.

Our main result is that in all string ensembles and random matrix theories that we study, $g_{a\gamma\gamma}$ increases

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

polynomially in N . We perform detailed studies of two string geometry ensembles that we refer to as the tree ensemble and the hypersurface ensemble, using the mass threshold $m_a \leq 10^{-12}$ eV so that our results can be compared to numerous experiments. In the tree ensemble at the expected value of $N = 2483$, we find that the mean of the projected $g_{a\gamma\gamma}$ distribution is $g_{a\gamma\gamma} = 3.2 \times 10^{-12}$ GeV $^{-1}$. In the hypersurface ensemble $N = 491$ is expected, and we find the expected value $g_{a\gamma\gamma} = 2.0 \times 10^{-10}$ GeV $^{-1}$. We also study the F-theory geometry with the most flux vacua [15], which is very constrained due to the existence of only a single divisor that can support the Standard Model, given our assumptions; it yields $g_{a\gamma\gamma} = 3.47 \times 10^{-12}$ GeV $^{-1}$. All of these couplings are in range of future experiments, and in some cases are already in tension with data; see the discussion. Furthermore, by removing the mass threshold the expected value of $g_{a\gamma\gamma}$ does not change significantly, implying that ALPs in string ensembles that we study will not be seen in searches at LHC or future colliders. The random matrix results suggest that $g_{a\gamma\gamma}$ should increase significantly with N , even if our vacuum does not arise in one of the studied ensembles.

We emphasize at the outset that we are *modeling* ALP-photon couplings using data from string theory. A complete calculation, with engineered Standard Models and full moduli stabilization, is computationally intractable given current techniques, which can be formalized in the language of computational complexity [16]. At small N , however, more complete calculations can be performed, including partial moduli stabilization; see, e.g., [17]. Instead, we model ALP-photon couplings by intelligently sampling the Calabi-Yau moduli space and utilizing knowledge of how realistic gauge sectors arise, without trying to concretely engineer them or stabilize moduli. This allows for the study of large ensembles at large N , but also motivates further studies once techniques for the complete calculations become available. The models that we use are described thoroughly in the text and are accurately summarized in the discussion.

This paper is organized as follows. In Sec. II we discuss ALP-photon couplings from the perspective of string theory. In Sec. III we introduce the ensembles of string compactifications and random matrix effective theories. In Sec. IV we present the $g_{a\gamma\gamma}$ distributions computed in the ensembles. We review constraints and projections of existing and proposed experiments in Sec. V and discuss our results in light of them in Sec. VI. Interesting future directions are also discussed.

II. ALP-PHOTON COUPLINGS IN STRING THEORY

Low-energy effective field theories from string theory regularly have a large number of ALPs, and under certain assumptions related to control of the theory many of them

are very light, as we will soon discuss. The Lagrangians of interest, focusing on the ALP sector, take the form

$$\mathcal{L} = -\frac{1}{4}F^{\mu\nu}F_{\mu\nu} - \frac{1}{2}\delta_{ij}(\partial^\mu\phi^i)(\partial_\mu\phi^j) - m_i^2(\phi^i)^2 - \frac{1}{4}c_i\phi^i\tilde{F}^{\mu\nu}F_{\mu\nu}. \quad (3)$$

Here $F_{\mu\nu}$ is the electromagnetic field strength, with $\tilde{F}^{\mu\nu} = \epsilon^{\mu\nu\gamma\sigma}F_{\gamma\sigma}$, and ϕ^i are the ALPs, with masses m_i . From a low-energy perspective these parameters are generally unconstrained; however, we will find that UV structure from string theory actually constrains the low-energy physics, and introduces correlations and patterns, that will result in interesting structure in both the ALP masses and the ALP-photon couplings.

For concreteness, we will study the dependence of ALP-photon couplings on the number of ALPs N in string compactifications and random matrix models. In the latter context, we consider compactifications of type IIB string theory/F-theory on a suitable Kähler manifold B , which yields a four-dimensional (4D) $\mathcal{N} = 1$ effective field theory (EFT). To date, such compactifications encompass the largest known portion of the $\mathcal{N} = 1$ landscape [15, 18–21]. Some of the most generic phenomena in this region, such as large gauge sectors and numbers of ALPs, correlate strongly with moving away from weakly coupled limits [22]. The ability to make such statements away from weak string coupling relies critically on the holomorphy implicit in algebraic geometry.

The data of such an F-theory compactification, in addition to the choice of B , are an elliptic fibration over B that encodes the data of the gauge group; see [23] for further details. The ALPs θ^i that we study arise from the dimensional reduction of the Ramond-Ramond four-form C_4 and become the imaginary parts of complexified Kähler moduli, written as

$$T^i = \int_{D_i} \left(\frac{1}{2}J \wedge J + iC_4 \right) \equiv \tau^i + i\theta^i. \quad (4)$$

Here the D_i are a basis of divisors (4-cycles) in B , numbering $h^{1,1}(B)$, and J is the Kähler form on B , which can be written in terms of parameters t_i as $J = t_i\omega^i$, with $\omega^i \in H^{1,1}(B)$. The τ^i parametrize the volumes of the divisors in B . A typical B has $h^{1,1}(B) \sim \mathcal{O}(10^3)$ [19, 24]; that is, in these ensembles, thousands of ALPs are expected.

In addition to the many ALPs, EFTs arising from string theory often contain many gauge sectors. Throughout we will take N to be the number of ALPs. In the context of F-theory, our focus will be on when this gauge sector, indexed by α , is supported on a stack of seven-branes wrapping a cycle Q^α , or from a nontrivial Mordell-Weil group element of the elliptic fibration Q^α . The ALP-gauge portion of the EFT takes the form

$$\mathcal{L} = -M_p^2 K_{ij}(\partial^\mu \theta^i)(\partial_\mu \theta^j) - V(\theta) - \sum_\alpha Q_\alpha^\alpha (\tau^i G_\alpha^{\mu\nu} G_{\alpha\mu\nu} + \theta^i \tilde{G}_\alpha^{\mu\nu} G_{\alpha\mu\nu}). \quad (5)$$

Here M_p is the Planck mass, $V(\theta)$ is the nonperturbative ALP potential, and K_{ij} is the metric on moduli space, which at tree level is derived from the Kähler potential $\mathcal{K} = -2 \log \mathcal{V}$, where $\mathcal{V}$ is the volume of B , which is expressed as

$$\mathcal{V} = \int_B J \wedge J \wedge J = \frac{1}{6} \kappa^{ijk} t_i t_j t_k, \quad (6)$$

and κ^{ijk} are the triple intersection numbers of B .

To justify our use of the tree-level Kähler potential, we note that F-theory can be viewed as the dimensional reduction of type IIB SUGRA in the presence of a spatially varying axio-dilaton. The type IIB SUGRA action is unique up to field redefinitions and higher-derivative corrections, and so any physical correction to the 4D EFT must come with additional powers of appropriate volumes of cycles. At large enough volume we therefore expect that using $\mathcal{K} = -2 \log \mathcal{V}$ should be a good approximation to the Kähler potential. It is not known how large is large enough in this case, but since the volumes in compactifications with many cycles tend to be very large [25], we believe this approximation should be valid for our study.

The leading-order Kähler potential is independent of the ALPs, so one can move to a canonically normalized frame in which the ALP mass matrix is diagonal. Let $F_{1\mu\nu} \equiv F_{\mu\nu}$ be the electromagnetic field strength, corresponding to a homology class Q (subtleties associated with electroweak symmetry breaking will be discussed momentarily). The Lagrangian of interest then takes the form

$$\mathcal{L} = -\frac{1}{2} \delta_{ij} (\partial^\mu \phi^i) (\partial_\mu \phi^j) - \frac{1}{4} F^{\mu\nu} F_{\mu\nu} - \frac{1}{4} c_i \phi^i \tilde{F}^{\mu\nu} F_{\mu\nu} - m_i^2 (\phi^i)^2. \quad (7)$$

In terms of geometric data, we can write [24]

$$|\tilde{c}| = \frac{\sqrt{Q \cdot K^{-1} \cdot Q}}{2M_p Q \cdot \tau}, \quad (8)$$

where K^{-1} is the inverse Kähler metric on field space. Equation (8) is independent of homogeneous scaling in the Kähler cone $J \rightarrow \lambda J$, and therefore only depends on the angle in the Kähler cone.

It is interesting to note that if $N = 1$, then this coupling is fixed to be $c = 1/(\sqrt{3}M_p)$ independent of the details of the geometry [26]. This is important because in the case of a single ALP it determines the photon coupling, $g_{a\gamma\gamma} \simeq 1/M_p$, which is well below current and (projected)

future experimental bounds. We will investigate whether large N effects lead to a significant enhancement of the coupling. Note that if the masses of all the ϕ^i that appeared nontrivially in $c_i \phi^i$ were the same, then $|c|$ would simply be the coupling of the canonically normalized field $\varphi \equiv c_i \phi^i / |c|$ to $F_{\mu\nu}$. In general this will not be the case; however, in string compactifications with N large we expect that some of the ALPs will be essentially massless, which we will now review.

ALP masses depend critically on the fact that, in the absence of sources, the θ^i enjoy a continuous shift symmetry to all orders in perturbation theory, broken to a discrete shift symmetry by nonperturbative effects. In particular, the superpotential W is known to receive corrections from stringy instantons and/or strong gauge dynamics that generate masses for the ALPs [27]. Due to the shift symmetry of the ALP and the holomorphy of the superpotential, any nonperturbative contribution to the superpotential takes the schematic form

$$\Delta W \sim e^{-2\pi \tilde{Q}_i (\tau^i + i\theta^i)}, \quad (9)$$

for rational $\tilde{Q}_i$. Therefore, for a given θ^i , if the corresponding τ^i is very large, then any nonperturbative contribution to W involving θ^i will be negligible, and θ^i is expected to be essentially massless (the same is not true of the τ^i , as they do not enjoy a similar shift symmetry and can receive masses via perturbative corrections to the Kähler potential [28,29]). The central observation of [25] is that in compactifications with large N , the region of moduli space where the EFT is expected to be valid (known as the stretched Kähler cone) is quite narrow. Restricting to the stretched Kähler cone for the sake of control forces some of the τ^i to be very large, which in turn forces some of the θ^i to be extremely light. This is the key result that (in this context) puts numerous ALPs in the sub-eV mass range relevant for the experiments we will discuss. In this work we make the technical assumption that the Standard Model sector does not generate a large mass term for the ALP ($\gtrsim 10^{-12}$ eV). This assumption is quite mild given the scales of the potential generated by $SU(2)_L$ instantons, as well as current expectations for the QCD axion mass [30–32].

If the ALP masses are much lower than the typical energies of an experiment, then they can be safely neglected and the ALPs taken to be massless. We will assume that the typical experiment energies are $\gtrsim 10^{-12}$ eV, and so any ALP with a mass $\ll 10^{-12}$ eV can be treated as massless. In this case we can define a single ALP $a \equiv c_i \phi^i / |c|$ that couples to $F_{\mu\nu}$, with strength $g_{a\gamma\gamma}$. Here the sum includes only the ϕ^i that have negligible mass. The relevant terms in Eq. (7) become

$$\mathcal{L} \supset -\frac{1}{2}(\partial^\mu a)(\partial_\mu a) - \frac{1}{4}F^{\mu\nu}F_{\mu\nu} - \frac{1}{4}g_{a\gamma\gamma}a\tilde{F}^{\mu\nu}F_{\mu\nu}, \quad (10)$$

where

$$g_{a\gamma\gamma} = \sqrt{\sum_i c_i^2}, \quad (11)$$

in terms of the couplings c_i in Eq. (7). Note that we can always redefine our ALP a to make $g_{a\gamma\gamma}$ non-negative, as we do henceforth.

It is useful to express Eq. (11) in terms of the original geometric quantities that appear in Eq. (5). To do so, we determine the massless axions by finding the linear combinations of axions that receive a non-negligible mass term (see below for a discussion) in the canonically normalized frame. Let the massive axions be specified by a (generally nonfull rank) matrix M_i^a , such that $\theta^a \equiv M_i^a \theta^i$ receives a non-negligible superpotential contribution in Eq. (9). We can move to a canonically normalized frame by writing $K = S^T \cdot f \cdot f \cdot S$, where S is a matrix of the orthonormal eigenvectors of K , and f is a diagonal matrix of the square roots of the eigenvalues. The massless axions (in the canonically normalized frame) are then specified by the matrix

$$D = \text{Ker}(M \cdot S^T \cdot f^{-1}); \quad (12)$$

i.e., the massless linear combination of canonically normalized axions is encoded in the rows or columns.

We can express these axions in the geometric θ basis by writing

$$Q^M = D \cdot f \cdot S, \quad (13)$$

where the superscript refers to “massless,” and we may therefore rewrite $g_{a\gamma\gamma}$ in Eq. (11) as

$$g_{a\gamma\gamma} = \frac{\sqrt{Q^M \cdot K^{-1} \cdot Q^M}}{2M_p Q \cdot \tau}, \quad (14)$$

for a fixed gauge group specified by the homology class Q .

Finally, we comment briefly on the string scale $M_s \simeq M_p/\sqrt{\mathcal{V}}$. For reasons that we will discuss, obtaining control over the string effective theory at large N requires going to regions in Kähler moduli space where a nontrivial number of four-cycle volumes is large. This correlates strongly with a large overall volume $\mathcal{V}$, which generally gives rise to an intermediate string scale at large N in our ensembles, $M_s \simeq 10^{12}\text{--}10^{15}$ GeV. However, since it is divisor volumes that more readily appear in our $g_{a\gamma\gamma}$ calculations and the geometry does not necessarily have the Swiss cheese property, we prefer to think of the relationship between M_s and $g_{a\gamma\gamma}$ as correlative, rather

than causal. It would be interesting to explore this further in future work.

Our goal is to compute $g_{a\gamma\gamma}$ in large ensembles of string compactifications, as well as random matrix models, in order to model the distribution of such couplings that one expects from string theory and to understand the behavior as the number of ALPs N grows large.

III. ENSEMBLES

Having reviewed the EFTs expected from string theory, the specific contexts in which we will study them, and the presence of ultralight ALPs, we now introduce the ensembles in which we perform this study.

A. The tree ensemble

F-theory is a nonperturbative generalization of type IIB string theory that allows for regions of strong string coupling and has a more general gauge spectrum than its IIB counterpart. In F-theory, the internal space B for compactification determines a minimal (geometric) gauge structure from non-Higgsable seven-branes, whose presence requires no tuning in complex structure moduli space.

In [19] a lower bound of the number of bases B suitable for F-theory compactifications was determined to be $4/3 \times 2.96 \times 10^{755}$ via the discovery of a construction algorithm for an ensemble of B known as the tree ensemble. This ensemble represents a large graph, where nodes are geometries and edges are simple topological transitions (known as blowups) that may take place between the geometries. More specifically, each node contains a parametric family (the complex structure moduli space) of geometries that have the same topological type, but for some representatives of the family the space becomes singular enough to allow for a topological transition. This means that at leading order in the physics, the space can be explored by movement along flat directions of the scalar potential to reach the singular point, from which a transition may be made to a different topological type (node) of geometry and accordingly a different EFT. Mathematically, this process may be done continuously in moduli space, as theorems relating canonical singularities and the Weil-Petersson metric ensure that the paths through the graph are at finite distance in moduli space. See [19,33–35] for further discussion.

The tree ensemble, while enormous in size, has certain tractable aspects, due to detailed understanding of the construction algorithm. In particular, the geometric gauge group can be determined to high accuracy. Such a set provides a rich ensemble in which to address questions about distributions of effective field theories in string theory. The tree ensemble is constructed by starting with a weak Fano toric variety B and performing blowups of B that satisfy sufficient conditions to remain at finite distance in moduli space. Such blowups can be performed over toric

points or curves, and each sequence of blowups of a particular toric point or curve in X is called a tree. A typical EFT from the tree ensemble has a minimal gauge group of the form

$$G \geq E_8^{10} \times F_4^{18} \times U^9 \times F_4^{H_2} \times G_2^{H_3} \times A_1^{H_4}, \quad (15)$$

where U is a B -dependent gauge group and H_i are computable B -dependent integers that are almost always nonzero. Other gauge groups can be tuned, but the group given in Eq. (15) is required by the geometry (though some of the gauge factors could be broken by the introduction of G_4 flux).

In addition, an overwhelming fraction of EFTs from the tree ensemble has a large number of ALPs. The expected number of ALPs is $N = 2483$. This value is determined by a bubble cosmology model on the tree ensemble, where the vacuum transitions are modeled by the topological transitions between the geometries [35] (drawing from a flat distribution gives similar results, with the preferred value of $N = 2015$).

From the expected gauge structure of the tree ensemble discussed in the previous section it is clear that there are in principle many ways to realize the Standard Model in the tree ensemble. In addition to the minimal gauge structure given in Eq. (15) one can tune additional gauge groups, and additional $U(1)$ factors could be realized by sections in the elliptic fibration.

For understanding the couplings of the ultralight ALPs to the photon, the relevant gauge sector is $SU(2)_L \times U(1)_Y$. Since pure Abelian factors have not been studied in the tree ensemble, we will consider two cases of gauge sectors. First, one could try to realize $SU(2)_L \times U(1)_Y$ by embedding it in a larger non-Abelian gauge group, for instance a GUT, in which case we will model the coupling $g_{a\gamma\gamma}$ directly as the coupling of the ALP to the larger group, assuming it breaks to the Standard Model in one of the canonical ways. The second case is realizing $SU(2)_L$ directly from the geometrically determined $SU(2)$ gauge symmetry on the seven-brane, in which case we will compute the contribution to $g_{a\gamma\gamma}$ from $SU(2)_L$. The cases are summarized as follows:

- (1) $G = SU(2)_L$ arises directly on a seven-brane.
- (2) G is a geometric gauge group on a seven-brane such that $G \rightarrow SU(2)_L \times U(1)_Y \times G'$, but $G' \not\supset SU(3)$; i.e., QCD comes from a different seven-brane.
- (3) G is a gauge group on seven-brane such that $G \rightarrow SU(2)_L \times U(1)_Y \times SU(3)$; this includes some common GUT scenarios.

In the last option the ALP linear combination that couples to $SU(2)_L$ is the QCD axion, while in the first two options this is not necessarily the case. Since all three possibilities lead to similar results, the forthcoming plots of ALP-photon couplings take into account all three.

In case 1, where we compute the contribution to $g_{a\gamma\gamma}$ from $SU(2)_L$, which we denote g_{aWW} , the coupling to the physical photon is computed as

$$g_{a\gamma\gamma} = g_{aWW} \sin^2 \theta_w + g_{aYY} \cos^2 \theta_w, \quad (16)$$

where θ_w is the Weinberg angle ($\sin^2 \theta_w \simeq 0.23$), and g_{aYY} is the coupling to $U(1)_Y$. In only this case it is possible that $g_{aYY} \neq g_{aWW}$. However, as modeling the hypercharge in our ensembles is nontrivial, we simply assume that the hypercharge contribution to $g_{a\gamma\gamma}$ does not lead to significant cancellation. This would require g_{aWW} and g_{aYY} to be of similar order of magnitude and opposite sign. As long as this is not the case, which we find plausible, $|g_{aWW}|$ gives an approximate lower bound to $|g_{a\gamma\gamma}|$.

B. Calabi-Yau hypersurfaces in toric varieties

In addition to the tree ensemble, we consider Calabi-Yau hypersurfaces in toric varieties [36,37]. Such a hypersurface combinatorially corresponds to a triangulated reflexive 4D polytope. Systematically orientifolding a large number of Calabi-Yau threefolds is beyond the scope of this work, and so we use the geometric data obtained from the Calabi-Yau itself as a model for the appropriate $\mathcal{N} = 1$ data, namely in constructing K^{-1} . For the hypersurfaces case, we assume that the gauge group (see the listed options in Sec. III A) is supported on the restriction of a divisor Q that is a linear combination of toric divisors to the hypersurfaces. There is, *a priori*, an infinite number of choices for such a linear combination, but physical considerations will render this set finite and computable, as we will discuss in Sec. IV.

C. Random matrix EFTs

For the sake of comparison to our string results, we will also compute ALP-photon couplings in certain random matrix (RM) effective field theories. We emphasize, though, that we do not currently have a reason to believe that the random matrix ensembles that we study accurately represent actual string data [38]. Instead, we simply wish to compare and also to demonstrate that they can give rise to growing ALP-photon couplings as a function of N . This lends some further credence to the idea that $g_{a\gamma\gamma}$ should increase with N .

For our random matrix theory (RMT) analysis we will consider the simplified Lagrangian of massless ALPs and gauge fields,

$$\mathcal{L} = -\frac{M_p^2}{2} K_{ij} (\partial^\mu \theta^i) (\partial_\mu \theta^j) - \frac{1}{4} F^{\mu\nu} F_{\mu\nu} - \frac{1}{4} Q_i \theta^i \tilde{F}^{\mu\nu} F_{\mu\nu}, \quad (17)$$

where $F_{\mu\nu}$ is the electromagnetic field strength. In a canonically normalized frame the ALP-photon coupling is simply

$$g_{a\gamma\gamma} = \sqrt{Q \cdot K^{-1} \cdot Q}. \quad (18)$$

A natural model for K (K^{-1}) is to draw it from an (inverse)-Wishart distribution, as we generally observe that the entries of K (K^{-1}) shrink (grow) as a function of N . Taking K to be a Wishart matrix, we have $K = A^\dagger A$, with the entries of A drawn from a normal distribution $\Omega(0, \sigma)$ centered around zero with standard deviation (SD) σ . This in turn makes K^{-1} an inverse-Wishart matrix, which in practice is easier to generate directly than to invert K . We take Q to be a unit vector and study two cases: the first being the (unnormalized) entries of Q drawn randomly from the distribution above, and the second taking Q to be a unit vector pointing in a basis direction $Q = \hat{e}_i$.

IV. DISTRIBUTIONS OF ALP-PHOTON COUPLINGS IN THE ENSEMBLES

We wish to understand the distribution of $g_{a\gamma\gamma}$ in the generic case, i.e., when N is large. Computing $g_{a\gamma\gamma}$ for the largest N regime of our ensembles in a large number of examples is computationally prohibitive, and so in order to explore the large N regime we will study $g_{a\gamma\gamma}$ for moderately large N , and then extrapolate to the largest N , where the bulk of the geometries are believed to occur (this has in fact been demonstrated explicitly in a cosmological model on the tree ensemble, and there is strong evidence for this in the case of hypersurfaces [39]). We will find that in fact the relevant statistical quantities in the distributions of $g_{a\gamma\gamma}$ obey nice scaling properties.

A. The EFT computation

1. Constructing the ensembles

To construct the ensemble of geometries to study we will randomly draw geometries from the various ensembles. For the tree ensemble, we construct 1000 geometries for one through ten trees each, corresponding to $h^{1,1}(B)$ ranging from 55 to 235, differing by jumps of $\Delta h^{1,1}(B) = 20$. For a fixed number of trees we randomly draw a tree configuration from the ensemble and blow up a randomly drawn toric point with that sequence of blowups. The configuration of blowups fixes the topological properties of B .

For the case of hypersurfaces, we randomly draw reflexive 4D polytopes from the Kreuzer-Skarke ensemble [37] and compute the pushing triangulation of each polytope to calculate the relevant topological data for the corresponding Calabi-Yau hypersurface. We select 1000 geometries for $h^{1,1}(B) = 10, 20, 30, 40, 50, 80, 120, 160$. For $h^{1,1} = 200$ there are only 706 polytopes; we utilized all of them.

Having calculated the relevant topological data, in order to write down the effective field theory for the ALPs and gauge sectors, we need to specify the VEVs of the τ^i . This is done by choosing a point in the Kähler cone.

2. The Kähler cone

In a complete calculation, the VEVs of the τ^i are determined by moduli stabilization. We make the assumption that the τ^i are stabilized in a regime of non-perturbative control. In particular, we require the volumes of all curves to be greater than or equal to unity. This region is known as the stretched Kähler cone, and as observed in [25], is very narrow at large $h^{1,1}(B)$. As in [24], where the tree ensemble was explored in the context of axion reheating, we evaluate the τ^i at the apex of the stretched Kähler cone, defined by minimizing the sum of the toric curves. Since the cone is narrow, and the couplings in Eq. (8) are invariant under scaling out in the cone via $J \rightarrow \lambda J$, we expect this point to be a good representative of the physics. In addition, as one scales out $J \rightarrow \lambda J$ the cycle volumes increase, and so more ALPs will become light, which would enhance the couplings of the gauge groups to the ultralight ALPs, and so performing this analysis should provide a lower bound on the size of the couplings in the stretched Kähler cone.

3. ALP masses and gauge couplings

With the EFT data in hand, we proceed to compute $g_{a\gamma\gamma}$. Since we are interested in the effects of the ultralight ALPs we only need to determine which ALPs will have masses much less than any relevant experimental energy, which we take to be 10^{-12} eV. Recall that for any ALP θ^i , any term in the potential is accompanied by an exponentially suppressed prefactor $\exp(-2\pi\tau^i)$, and therefore the θ^i who have $\tau^i \gg 1$ will have very small masses. In the full $\mathcal{N} = 1$ SUGRA potential such terms are accompanied by prefactors involving inverse powers of $\mathcal{V}$, the constant term W_0 arising from the Gukov-Vafa-Witten (GVW) flux superpotential [40], and the various two- and four-cycle volumes, and such prefactors are generally $\ll \mathcal{O}(1)$ in the stretched Kähler cone (see [25]). In general, linear combinations of divisors can contribute to the superpotential, so to give an upper bound for the ALP mass one should consider a generating set of divisors such that all effective divisors can be expressed as non-negative integer linear combinations of that set. For a toric variety the toric divisors generate the cone of effective divisors, known as the effective cone.¹ We can therefore give an upper bound on the mass scale generated for an ALP θ^i of the form

$$m_i^2 \lesssim \frac{1}{f_{\min}^2} e^{-2\pi\tau^i/C_i}, \quad (19)$$

where $f_{\min}$ is the smallest ALP decay constant (square root of the smallest eigenvalue of K), the τ^i are the volumes of

¹In the hypersurface case we work under the assumption that the effective cone of the ambient space provides a good approximation to the effective cone of the hypersurface.

the generators of the effective cone, and C_i are geometry dependent constants that range from 1 to 30 and are often dual Coxeter numbers. The inverse factor of $f_{\min}^2$ provides the weakest upper bound of the canonical normalization effects. In practice, many θ^i have their corresponding $\tau^i \gg 1$, and so we can treat such ALPs as massless. Concretely, this means that we compute the upper bounds on the masses using (19) and the values of $f_{\min}$ and τ_i in a geometry, and call the ALP θ_i massless if the associated bound is below a fixed mass threshold. For instance, in the tree ensemble the fraction of axions with masses $\leq 10^{-12}$ eV grows from 0.37 at $N = 55$ to 0.46 at $N = 195$, at which it becomes approximately constant.

Keeping only these nearly massless ALPs, i.e., those with a mass upper bound below 10^{-12} eV, we canonically normalize the ALPs and gauge fields. Not every gauge group is a viable candidate for the SM as mentioned in Sec. III A. In particular, the gauge coupling in the UV must be large enough to produce the correct low-energy gauge couplings. For a gauge group supported on a cycle given by Q_i , we have the relation

$$\frac{1}{g_{UV}^2} \simeq Q_i \tau^i, \quad (20)$$

and so any cycle with $Q_i \tau^i$ very large gives rise to very weak gauge couplings that are not consistent with well-studied models and the observed gauge couplings. For instance, the SUSY GUT value is $\alpha_{UV} \simeq 0.03$, which corresponds to $Q_i \tau^i = 2.7$. Leaving some room for model building, we demand that α_{UV} is not more than 1 order of magnitude smaller than the SUSY GUT value, and therefore we impose the cutoff $Q_i \tau^i \leq 25$. In particular, in the hypersurface ensemble we study SM candidates on linear combinations Q_i of toric divisors whose volume satisfies $Q_i \tau^i \leq 25$, while in the tree ensemble we study SM candidates on toric divisors with non-Higgsable clusters with volume ≤ 25 .

For each gauge group satisfying this condition, we then compute the coupling of the corresponding ultralight ALP. We present the results in the next section.

4. ALP decays to other sectors

A generic ALP in our setup couples to a large number of gauge sectors (other than the visible sector) via terms of the form $c_{i,\alpha} \phi^i \tilde{G}_\alpha^{\mu\nu} G_{\alpha\mu\nu}$ for fixed α (note that this coupling is written for canonically normalized fields). An experimental concern arises if the ALP decays too quickly to another sector; for instance, in helioscope experiments the ALPs may in principle decay into a dark sector before it travels from the sun to Earth. The decay rate of an ALP of mass m_i to a pair of dark gluons in one of these sectors is given by

$$\Gamma_{\phi^i \rightarrow gg} = \dim(G) \frac{c_{i,\alpha}^2 m_i^3}{64\pi}. \quad (21)$$

The coupling may be written

$$c_{i,\alpha} \approx (16.6 \text{ GeV}^{-1}) \times \dim(G)^{-1/2} \left(\frac{1 \text{ eV}}{m_i} \right)^{3/2} \left(\frac{8 \text{ min}}{\Gamma_{\phi^i \rightarrow gg}^{-1}} \right)^{1/2}. \quad (22)$$

Therefore, the coupling has to be $\mathcal{O}(10)$ GeV^{-1} for an eV ALP produced in the Sun to decay before it reaches the Earth. In all of our ensembles, the couplings of ALPs to any of the gauge sectors are orders of magnitude less than 10^{-1} GeV^{-1} , and therefore premature decays of ALPs will not be a concern. The ALPs will not decay before reaching the Earth.

B. Results

1. Results for the tree ensemble

In Fig. 1 we show the normalized distributions of $\log_{10}(g_{a\gamma\gamma} \times \text{GeV})$ for our smallest $N = 55$ and largest $N = 235$. Clearly the distribution shifts to the right as N grows.

In order to extrapolate to even larger values of N , which is often argued to be the location of the most vacua in the string landscape, we will determine quantitative properties of the N dependence of the distribution. In particular, we find that both the mean and the SD of $\log_{10}(g_{a\gamma\gamma} \times \text{GeV})$ have well-behaved N scaling, as shown in Fig. 2, and are therefore suitable to use for such an extrapolation.

We fit a power law to the mean of $\log_{10}(g_{a\gamma\gamma} \times \text{GeV})$ as a function of N . Expressed in terms of $g_{a\gamma\gamma}$, we find




FIG. 1. The normalized distributions of $\log_{10}(g_{a\gamma\gamma} \times \text{GeV})$ for our smallest $N = 55$ and largest $N = 235$, as well as the extrapolated distribution for the preferred value of $N = 2483$ in the tree ensemble. There is a clear shift of the distribution towards larger values as N grows. Current (solid) and projected (dashed) exclusion lines are presented for various experiments.




FIG. 2. Top: The mean of $\log_{10}(g_{\gamma\gamma} \times \text{GeV})$ in the tree ensemble as a function of N , with a best-fit curve. Bottom: The SD of $\log_{10}(g_{\gamma\gamma} \times \text{GeV})$ in the tree ensemble as a function of N , with a best-fit curve.

$$\text{mean}(g_{\gamma\gamma}) = 2.73 \times 10^{-18} \times N^{1.77} \text{ GeV}^{-1}, \quad (23)$$

which gives the dependence of the coupling on the number of ALPs. We note here that “mean” indicates that we have fit the mean of $\log_{10}(g_{\gamma\gamma} \times \text{GeV})$, as opposed to $g_{\gamma\gamma}$ itself. Given the excellent fit, it is reasonable to extrapolate the mean of the distribution of $\log_{10}(g_{\gamma\gamma} \times \text{GeV})$ to the preferred value of N in the tree ensemble, at $N = 2483$. At this value the predicted mean of $\log_{10}(g_{\gamma\gamma} \times \text{GeV})$ is -11.50 , which is much larger than the mean at small-to-moderate N . In order to estimate the distribution itself at $N = 2483$ we assume that the distribution is modeled by a Gaussian, with the mean and standard deviation obtained by extrapolating the curves in Fig. 2 to $N = 2483$. The expected distribution is shown in Fig. 1, along with its smaller- N counterparts, for the sake of comparison.

For the tree ensemble at large N , the distribution sits right on the edge of experimental sensitivity, depending slightly on the assumptions and experiment. We will present a thorough analysis of our results relative to experimental results and prospects in the discussion, since it will be useful to also compare results across ensembles.

We would like to understand the origin of this result. It is clear from our analysis so far, and from Fig. 1, that as N increases so does the mean of $\log_{10}(g_{\gamma\gamma} \times \text{GeV})$. Recall




FIG. 3. A plot of $\text{mean}(\log_{10}(g_{\gamma\gamma} \times \text{GeV}))$ versus $\text{mean}(\log_{10}(\lambda_{\max}(K^{-1})))$. The slope of the line is ~ 0.38 .

from Eq. (8) that $g_{\gamma\gamma}$ is proportional to the norm of a vector Q^i , computed with K^{-1} . Since we are imposing $Q \cdot \tau \lesssim 25$ so that the cycle in question can support realistic SM gauge couplings, the denominator of Eq. (14) is essentially a constant, ranging from 1 to 25, compared to the large hierarchies that may arise in the numerator. In Fig. 3, we show a plot of $\text{mean}(\log_{10}(g_{\gamma\gamma} \times \text{GeV}))$ versus $\text{mean}(\log_{10}(\lambda_{\max}(K^{-1})))$, where $\lambda_{\max}(K^{-1})$ is the largest eigenvalue of K^{-1} . Figure 3 shows a clear correlation, which can be explained as follows: if the largest eigenvalue of K^{-1} grows with N , and if Q^i has nontrivial overlap with the corresponding eigenvector that does not shrink too quickly with N , then we will find N -dependent growth of $g_{\gamma\gamma}$ with N . This is indeed the case. In Fig. 4 we show the N dependence of $\text{mean}(\log_{10}(\lambda_{\max}(K^{-1})))$ with N . Clearly $\lambda_{\max}(K^{-1})$ grows rapidly with N . In addition, we find that the alignment of the Q^i with the corresponding eigenvector $\hat{v}_{\max}$ shrinks slowly as a function of N : at $N = 55$ we find $\text{mean}(\log_{10}(\hat{v}_{\max} \cdot \hat{Q})) = -1.72$, while at $N = 235$ we find $\text{mean}(\log_{10}(\hat{v}_{\max} \cdot \hat{Q})) = -2.59$. Therefore, the growth of $\lambda_{\max}(K^{-1})$ dominates over the decreasing $\hat{v}_{\max} \cdot \hat{Q}$, explaining the observed correlation.




FIG. 4. A plot of $\text{mean}(\log_{10}(\lambda_{\max}(K^{-1})))$ versus N , with a best-fit line to demonstrate the correlation.

From a random matrix perspective, constructing a RMT ensemble whose largest eigenvalue reproduces the scaling in Fig. 4 seems reasonable, as one could choose a Wishart ensemble whose distribution of entries reproduced such a scaling. However, given that $\hat{Q}$ tend to point along the standard basis $\hat{e}$ directions in the natural geometric basis, it is clear that the corresponding eigenvector $\hat{v}_{\max}$ does not obey standard eigenvector delocalization, as we would expect the entries of $\hat{v}_{\max}$ to be roughly $1/\sqrt{N}$.

2. Results for hypersurfaces

We perform the same analysis for hypersurfaces and find quite similar results. In particular, we find that the mean of $\log_{10}(g_{a\gamma\gamma} \times \text{GeV})$ increases as a function of N , even more rapidly than in the tree ensemble, and such an increase can be correlated with the maximal eigenvalue of K^{-1} .

In Fig. 5 we show the distributions for our smallest and largest N that we analyze in the hypersurface case, which are $N = 10$ and $N = 200$, respectively. In the same figure we also show the projected distribution for the largest N in the hypersurface ensemble, which is $N = 491$. There is a striking feature in N dependence of the distribution in the hypersurface case as compared to the tree ensemble: while the standard deviation of the distribution of $\log_{10}(g_{a\gamma\gamma})$ decreased as a function of N for the tree ensemble, in the case of hypersurfaces it actually increases, leading to the largest spread for large N . This could be due to the rich intersection structure of Calabi-Yau threefolds as compared to toric threefolds, whose intersection numbers are relatively tame. We fit a power law to the mean of $\log_{10}(g_{a\gamma\gamma} \times \text{GeV})$ as a function of N . Expressed in terms of $g_{a\gamma\gamma}$, we find

$$\text{mean}(g_{a\gamma\gamma}) = 8.52 \times 10^{-22} \times N^{4.22} \text{ GeV}^{-1}. \quad (24)$$




FIG. 5. The normalized distributions of $\log_{10}(g_{a\gamma\gamma} \times \text{GeV})$ for the hypersurfaces. Shown are the calculated distributions for $N = 10$ and $N = 200$, as well as the extrapolated distribution for the largest hypersurfaces at $N = 491$. There is a clear shift of the distribution towards larger values as N grows. Current (solid) and projected (dashed) exclusion lines are presented for various experiments.

For the hypersurface at large $N = 491$, the mean of the projected distribution is at $\log_{10}(g_{a\gamma\gamma} \times \text{GeV}) = -9.71$, and so a significant portion of the distribution is already in tension with data; see the discussion below.

3. Results for RMT

As described in Sec. III C we model K^{-1} as an inverse Wishart matrix (or K as a Wishart matrix) drawn from $\Omega(0, \sigma)$, for some σ , and Q a unit vector either pointing in a basis direction or with entries drawn from $\Omega(0, \sigma)$. We find that the different models for Q produce essentially the same behavior, so we will focus on the case that $Q = (1, 0, \dots, 0)$. In the simplified model the coupling of ALP to the photon is given by Eq. (18).

In Fig. 7 we show $\log_{10}(g_{a\gamma\gamma} \times \text{GeV})$ computed in the RMT models versus $\text{mean}(\log_{10}(\lambda_{\max}(K^{-1})))$, to demonstrate the similar behavior to the string ensembles studied above. Implicitly in the plot $\text{mean}(\log_{10}(\lambda_{\max}(K^{-1})))$ is increasing with N , and each data point is at a different N , ranging from 0 to 1500. Comparing the growth of $\text{mean}(\log_{10}(g_{a\gamma\gamma} \times \text{GeV}))$ with respect to both N and the maximum eigenvalue of K^{-1} , we note that the inverse Wishart ensembles with $\sigma = 1/N^2$ and $\sigma = 1/N$ are the closest fits to the actual string data, though both differ from it nontrivially.

This lends further credence to the central idea of this work: $g_{a\gamma\gamma}$ should grow significantly with N . In particular, if some other string ensemble has a different ALP-photon coupling scaling compared to the ones that we have studied, these EFT considerations suggest that the coupling should nevertheless increase polynomially in N , as shown in Fig. 6. We therefore believe that the central result likely extends to other string ensembles.




FIG. 6. The mean of $\log_{10}(g_{a\gamma\gamma} \times \text{GeV})$ in the inverse Wishart model for $\sigma = 1, 1/\sqrt{N}, 1/N, 1/N^2, 1/N^3$ as a function of N , with a best-fit line of the form $\text{mean}(g_{a\gamma\gamma}) \propto N^x$. The models with $\sigma = 1/N^2$ and $\sigma = 1/N$ are the best for the string data.




FIG. 7. $\log_{10}(g_{\gamma\gamma} \times \text{GeV})$, computed in the RMT models, versus $\text{mean}(\log_{10}(\lambda_{\max}(K^{-1})))$, with varying σ . Note that for all models both $\log_{10}(g_{\gamma\gamma} \times \text{GeV})$ and $\text{mean}(\log_{10}(\lambda_{\max}(K^{-1})))$ increase as a function of N .

4. Remarks on other ensembles

While the tree ensemble of F-theory bases and the Kreuzer-Skarke (KS) ensemble of Calabi-Yau (CY) hyper-surfaces are the largest-to-date explicit ensembles of string geometries, there are other important examples that we now discuss.

The first is that of the skeleton ensemble [21], which generalizes the tree ensemble beyond the simple sufficient criteria to remain at finite distance in moduli space. Actually, in this case the construction of the geometries B is a less-restrictive version of the tree ensemble, and so we expect the results to exhibit similar scaling. However, the largest number N of ALPs found in this ensemble is 16103, which is a great deal larger than the maximal N found in the tree ensemble. While this result should be taken with a grain of salt, since we do not perform the




FIG. 8. The normalized predicted distribution of $\log_{10}(g_{\gamma\gamma} \times \text{GeV})$ for the largest known value of $N = 16103$ in the skeleton ensemble, contrasted with that predicted for the largest $N = 2483$ in the tree ensemble. Current (solid) and projected (dashed) exclusion lines are presented for various experiments.

computation for reasons of complexity, it is interesting to check what the distribution of $g_{\gamma\gamma}$ would be if we were able to extrapolate outside of the tree ensemble to such an boundary point of the skeleton ensemble. We extrapolate using the fit of the related tree ensemble in Fig. 8.

Another important example is the geometry $B_{\max}$, which is the F-theory base thought to house the largest number of flux vacua [15]. In this geometry the number of ALPs is 98, and the non-Higgsable gauge group is $E_8^9 \times F_4^8 \times (G_2 \times SU(2))^{16}$. At the apex of the stretched Kähler cone there is only a single cycle with volume τ that supports a gauge group and satisfies $\tau \leq 25$. This cycle supports an E_8 group that could provide a GUT after flux breaking. For this gauge group we find $g_{\gamma\gamma} = 3.47 \times 10^{-12} \text{ GeV}^{-1}$. This coupling is also projected to be probed at future experiments.

V. CURRENT AND FUTURE EXPERIMENTS

Even though they are not relevant for our analysis, for the sake of thoroughness we will begin by discussing experiments that assume that ALPs contribute nontrivially to the dark matter and/or are sensitive only to ALPs in some finite mass region. Currently, the best limits on the existence of light ($m_a \lesssim 1 \text{ eV}$) axionlike particles comes from resonant-cavity haloscope experiments. The bounds are restricted to narrow mass windows centered around $m_a \simeq 2.5 \times 10^{-6} \text{ eV}$ for ADMX [41] and $m_a \simeq 2.5 \times 10^{-5} \text{ eV}$ for HAYSTAC [42]. Both experiments place an impressive limit on the coupling of $g_{\gamma\gamma} \lesssim 10^{-15} \text{ GeV}^{-1}$. However, the effectiveness of these searches is highly contingent on the assumption that the ALP in question comprises the totality of the local dark matter halo, as the signal scales linearly with the local halo density of ALPs. String theory generally provides a wealth of potential dark matter candidates, and even if an ALP were to represent the bulk of the local dark matter density, it would be a remarkable coincidence if this particle was to also couple to photons in an appreciable manner.

Considering the above shortcoming, we are thus led to consider experiments which are both broadband in sensitivity and independent of assumptions about the ALP contribution to local dark matter. The so-called “light-shining-through-walls” (LSW) experiments achieve both, with the DESY-based ALPS [43] and the CERN-based OSQAR [44] achieving a limit of $g_{\gamma\gamma} \lesssim 6 \times 10^{-8} \text{ GeV}^{-1}$ for ALPs with vanishing masses up to a threshold of $m_a \simeq 10^{-4} \text{ eV}$. The proposed ALPS-II experiment [45] hopes to achieve a sensitivity to couplings of order $g_{\gamma\gamma} \lesssim 2 \times 10^{-11} \text{ GeV}^{-1}$ in the very near future. Such a limit will be competitive with the current bound set by the CERN helioscope CAST, which constrains the coupling to $g_{\gamma\gamma} \lesssim 7 \times 10^{-11} \text{ GeV}^{-1}$ for ALPs with vanishing masses up to a threshold of $m_a \simeq 10^{-2} \text{ eV}$. This is the current best limit for ALPs in the broad mass range of interest to this paper.

Both of the techniques mentioned above should see improvements in future years. The International Axion Observatory (IAXO), a next-generation helioscope, should reach $g_{a\gamma\gamma} = 2 \times 10^{-12} \text{ GeV}^{-1}$ [3], with a prototype taking data within the next few years, while next-generation superconducting radiofrequency cavities hope to reach a similar limit on the same timescale [4].

Finally, we remark on the ability of astrophysical observations to constrain ALP couplings to photons. Oscillations in the spectrum of observed x-ray and gamma-ray photons from extragalactic sources may be induced by photon-ALP conversion in the presence of large galactic magnetic fields. This technique can be applied for very low ALP masses ($m_a \lesssim 10^{-9} \text{ eV}$), with the best limits coming from x-ray astronomy for $m_a \lesssim 10^{-12} \text{ eV}$. The most stringent current limits of $g_{a\gamma\gamma} \lesssim 8 \times 10^{-13} \text{ GeV}^{-1}$ come from observations of a number of point sources by the Chandra satellite [10]. In the next decade, both the European Space Agency (ATHENA) [11] and NASA (STROBE-X) [14] have proposed missions to improve these limits by an additional order of magnitude.

VI. DISCUSSION

In this paper we studied distributions of ALP-photon couplings in ensembles of string compactifications and random matrix theories. Since hundreds or thousands of axionlike particles typically arise in string compactifications, we studied the scaling of the effective $g_{a\gamma\gamma}$ distribution with N and its extrapolation to the expected value of N in a given ensemble.

Before stating the main results, we review the setup. The details of string geometry, which in our cases are Calabi-Yau manifolds in IIB compactifications or Kähler threefold bases of F-theory compactifications, determine the coupling of a linear combination of ALPs to gauge sectors arising on a four-manifold parametrized by an integral vector Q . Specifically, the string data that enter the calculation are the choice of manifold from the ensemble, taken to be at the apex of the stretched Kähler cone, the computation of the Kähler metric K_{ij} on ALP kinetic terms that plays a crucial physical role in canonically normalized couplings, and choices of Q that in principle allow for the Standard Model gauge couplings. ALP masses are generated by instanton corrections to the superpotential, and we utilize a natural geometric procedure that gives an upper bound on the ALP mass. The effective ALP-photon couplings that we study include only the contributions from individual ALPs whose mass upper bound is below an experimentally relevant threshold 10^{-12} eV . This conservative choice allows our calculations to also apply to experiments that are sensitive to ALPs below masses of 10^{-6} eV (e.g., [4]) and 10^{-2} eV (e.g., helioscopes). Further details are, of course, provided in the main text.

Our main results in each ensemble are that the mean value of the $g_{a\gamma\gamma}$ distribution ($\overline{g_{a\gamma\gamma}}$) scales polynomially in N , and its extrapolation to large N , where most vacua are argued to occur, is near the sensitivity of current or proposed experiments. We take each point in turn.

In both string ensembles we studied, the scaling $\overline{g_{a\gamma\gamma}}$ was significantly stronger than the $\sqrt{N}$ scaling that one might naively expect from Eq. (11). This is due to at least two effects. First, the number of ALPs contributing to the effective $g_{a\gamma\gamma}$ depends not only on N , but also the fraction of ALPs with masses below the chosen threshold; the latter itself increases with N [25]. Second, due to canonical normalization, $g_{a\gamma\gamma}$ depends critically on the largest eigenvalue of K^{-1} , which grows significantly with N ; see Figs. 3 and 4. Specifically, the two string ensembles that we studied were called the tree ensemble and the hypersurface ensemble, and they exhibited $\overline{g_{a\gamma\gamma}} \propto N^{1.77}$ and $\overline{g_{a\gamma\gamma}} \propto N^{4.22}$, respectively. We note that $\overline{g_{a\gamma\gamma}}$ can also be fit accurately to a weak exponential, but in the absence of any reason to expect exponential growth, we opt for the more conservative choice.

In both string ensembles that we studied we found that $\overline{g_{a\gamma\gamma}}$, extrapolated to the expected large value of N in each ensemble, is within range of current and/or proposed experiments. Specifically, with contributing ALPs having $m_a \leq 10^{-12} \text{ eV}$,

$$\begin{aligned} \text{tree: } \overline{g_{a\gamma\gamma}} &= 3.2 \times 10^{-12} \text{ GeV}^{-1} \quad \text{at } N = 2483, \\ \text{hyper: } \overline{g_{a\gamma\gamma}} &= 2.0 \times 10^{-10} \text{ GeV}^{-1} \quad \text{at } N = 491. \end{aligned} \quad (25)$$

We also computed $g_{a\gamma\gamma}$ in the F-theory geometry with the most flux vacua, where the requirement $\tau \leq 25$ within the stretched Kähler cone allows the SM to exist only on a single divisor, which carries geometric gauge group E_8 . In that example, $g_{a\gamma\gamma} = 3.47 \times 10^{-12} \text{ GeV}^{-1}$.

These results should be compared to the strongest current and projected constraints from experiments consistent with the assumptions of our analysis: that the ALPs do not have to (but could) comprise the dark matter, and that the masses are below a threshold rather than in a fixed window. For experiments satisfying these assumptions, the strongest current constraints are

$$\begin{aligned} \text{CAST: } g_{a\gamma\gamma} &\leq 7.0 \times 10^{-11} \text{ GeV}^{-1} \quad \text{for } m_a \leq 10^{-2} \text{ eV}, \\ \text{Chandra: } g_{a\gamma\gamma} &\leq 8.0 \times 10^{-13} \text{ GeV}^{-1} \quad \text{for } m_a \leq 10^{-12} \text{ eV}, \end{aligned} \quad (26)$$

and the strongest projected sensitivities in proposed experiments are

$$\begin{aligned} \text{IAXO: } g_{a\gamma\gamma} &\leq 2 \times 10^{-12} \text{ GeV}^{-1} \quad \text{for } m_a \leq 10^{-2} \text{ eV}, \\ \text{STROBE-X: } g_{a\gamma\gamma} &\leq 8 \times 10^{-14} \text{ GeV}^{-1} \quad \text{for } m_a \leq 10^{-12} \text{ eV}. \end{aligned} \quad (27)$$

We note the proximity of the computed $\overline{g_{\gamma\gamma}}$ in models of ALP-photon couplings in string compactifications to the experimental sensitivities.

Our study provides the first concrete evidence that ultralight ALPs in string compactifications could very well have detectable interactions with photons at large N .

Given this strong claim, we would like to clearly emphasize three potential caveats. First, this is a large N effect, and though there are compelling arguments that most vacua exist at large N , it could be that our vacuum is realized in the small N regime. Second, the quoted $\overline{g_{\gamma\gamma}}$ values are the mean of a projected distribution, and the photon associated with our vacuum could be realized in the tail. Third, the precise details of the distribution could be modified in more precise studies that attempt to explicitly realize Standard Model sectors and stabilize moduli, rather than modeling them, as we have. Each of these caveats is deserving of careful study.

Nevertheless, this result, taken together with the ubiquity of ALPs in string compactifications and the necessity of the photon in the correct theory of quantum gravity, provides excellent motivation for future work.

ACKNOWLEDGMENTS

We thank Joe Conlon, Caterina Doglioni, Yoni Kahn, Sven Krippendorff, Liam McAllister, Kerstin Perez, and Ben Safdi for valuable discussions. Portions of this work were completed at the Aspen Center for Physics, which is supported by National Science Foundation Grant No. PHY-1607611. J.H. is supported by NSF Grant No. PHY-1620526 and NSF CAREER Grant No. PHY-1848089. B.N. and G.S. are supported by National Science Foundation Grant No. PHY-1913328.

-
- [1] C. A. S. T. Collaboration, *Nat. Phys.* **13**, 584 (2017).
 - [2] C. P. Salemi (ABRACADABRA Collaboration), in *54th Rencontres de Moriond on Electroweak Interactions and Unified Theories (Moriond EW 2019) La Thuile, Italy* (2019), [arXiv:1905.06882](#).
 - [3] E. Armengaud *et al.* (IAXO Collaboration), *J. Cosmol. Astropart. Phys.* **06** (2019) 047.
 - [4] Z. Bogorad, A. Hook, Y. Kahn, and Y. Soreq, *Phys. Rev. Lett.* **123**, 021801 (2019).
 - [5] D. Wouters and P. Brun, *Astrophys. J.* **772**, 44 (2013).
 - [6] M. Berg, J. P. Conlon, F. Day, N. Jennings, S. Krippendorff, A. J. Powell, and M. Rummel, *Astrophys. J.* **847**, 101 (2017).
 - [7] M. C. D. Marsh, H. R. Russell, A. C. Fabian, B. P. McNamara, P. Nulsen, and C. S. Reynolds, *J. Cosmol. Astropart. Phys.* **12** (2017) 036.
 - [8] J. P. Conlon, F. Day, N. Jennings, S. Krippendorff, and M. Rummel, *J. Cosmol. Astropart. Phys.* **07** (2017) 005.
 - [9] F. Day and S. Krippendorff, [arXiv:1907.07642](#).
 - [10] C. S. Reynolds, M. C. D. Marsh, H. R. Russell, A. C. Fabian, R. Smith, F. Tombesi, and S. Veilleux, [arXiv:1907.05475](#).
 - [11] J. P. Conlon, F. Day, N. Jennings, S. Krippendorff, and F. Muia, *Mon. Not. R. Astron. Soc.* **473**, 4932 (2018).
 - [12] K. R. Dienes and B. Thomas, *Phys. Rev. D* **85**, 083523 (2012).
 - [13] J. Halverson, B. D. Nelson, and F. Ruehle, *Phys. Rev. D* **95**, 043527 (2017).
 - [14] P. S. Ray *et al.* (STROBE-X Science Working Group), [arXiv:1903.03035](#).
 - [15] W. Taylor and Y.-N. Wang, *J. High Energy Phys.* **12** (2015) 21.
 - [16] J. Halverson and F. Ruehle, *Phys. Rev. D* **99**, 046015 (2019).
 - [17] M. Cicoli, M. Goodsell, and A. Ringwald, *J. High Energy Phys.* **10** (2012) 146.
 - [18] W. Taylor and Y.-N. Wang, *J. High Energy Phys.* **01** (2016) 137.
 - [19] J. Halverson, C. Long, and B. Sung, *Phys. Rev. D* **96**, 126006 (2017).
 - [20] J. Halverson, C. Long, and B. Sung, *J. High Energy Phys.* **02** (2018) 113.
 - [21] W. Taylor and Y.-N. Wang, *J. High Energy Phys.* **01** (2018) 111.
 - [22] J. Halverson, *Nucl. Phys.* **B919**, 267 (2017).
 - [23] T. Weigand, *Proc. Sci., TASI2017* (2018) 016 [[arXiv:1806.01854](#)].
 - [24] J. Halverson, C. Long, B. Nelson, and G. Salinas, *Phys. Rev. D* **99**, 086014 (2019).
 - [25] M. Demirtas, C. Long, L. McAllister, and M. Stillman, [arXiv:1808.01282](#).
 - [26] M. Demirtas, C. Long, L. McAllister, and M. Stillman, [arXiv:1906.08262](#).
 - [27] E. Witten, *Nucl. Phys.* **B474**, 343 (1996).
 - [28] M. Cicoli, J. P. Conlon, and F. Quevedo, *J. High Energy Phys.* **10** (2008) 105.
 - [29] M. Cicoli, D. Ciupke, S. de Alwis, and F. Muia, *J. High Energy Phys.* **09** (2016) 026.
 - [30] G. G. di Cortona, E. Hardy, J. P. Vega, and G. Villadoro, *J. High Energy Phys.* **01** (2016) 034.
 - [31] L. McLerran, R. Pisarski, and V. Skokov, *Phys. Lett. B* **713**, 301 (2012).
 - [32] L. McLerran, *J. Exp. Theor. Phys.* **120**, 376 (2015).
 - [33] Y. Hayakawa, Degeneration of Calabi-Yau manifold with Weil-Petersson metric, Ph.D. thesis, University of Maryland, 1994, p. 44.
 - [34] C.-L. Wang, *Math. Res. Lett.* **4**, 157 (1997).
 - [35] J. Carifio, W. J. Cunningham, J. Halverson, D. Krioukov, C. Long, and B. D. Nelson, *Phys. Rev. Lett.* **121**, 101602 (2018).

- [36] V. V. Batyrev, [arXiv:alg-geom/9310003](#).
- [37] M. Kreuzer and H. Skarke, [Adv. Theor. Math. Phys.](#) **4**, 1209 (2000).
- [38] C. Long, L. McAllister, and P. McGuirk, [J. High Energy Phys.](#) **10** (2014) 187.
- [39] R. Altman, J. Carifio, J. Halverson, and B.D. Nelson, [J. High Energy Phys.](#) **03** (2019) 186.
- [40] S. Gukov, C. Vafa, and E. Witten, [Nucl. Phys.](#) **B584**, 69 (2000); **608**, 477(E) (2001).
- [41] N. Du *et al.* (ADMX Collaboration), [Phys. Rev. Lett.](#) **120**, 151301 (2018).
- [42] L. Zhong *et al.* (HAYSTAC Collaboration), [Phys. Rev. D](#) **97**, 092001 (2018).
- [43] K. Ehret *et al.*, [Phys. Lett. B](#) **689**, 149 (2010).
- [44] R. Ballou *et al.*, in *Proceedings, 10th Patras Workshop on Axions, WIMPs and WISPs (AXION-WIMP 2014): Geneva, Switzerland* (2014), pp. 125–130, [arXiv:1410.2566](#).
- [45] R. Bähre *et al.*, [J. Instrum.](#) **8**, T09001 (2013).