

PQ Axiverse

Mehmet Demirtas^{a,b}, Naomi Gendler^c, Cody Long^d, Liam McAllister^c, and Jakob Moritz^c

^a *The NSF AI Institute for Artificial Intelligence and Fundamental Interactions*

^b *Department of Physics, Northeastern University, Boston, MA 02115, USA*

^c *Department of Physics, Cornell University, Ithaca, NY 14853, USA and*

^d *Department of Physics and CMSA, Harvard University, Cambridge, MA 02138, USA*

We show that the strong CP problem is solved in a large class of compactifications of string theory. The Peccei-Quinn mechanism solves the strong CP problem if the CP-breaking effects of the ultraviolet completion of gravity and of QCD are small compared to the CP-preserving axion potential generated by low-energy QCD instantons. We characterize both classes of effects. To understand quantum gravitational effects, we consider an ensemble of flux compactifications of type IIB string theory on orientifolds of Calabi-Yau hypersurfaces in the geometric regime, taking a simple model of QCD on D7-branes. We show that the D-brane instanton contribution to the neutron electric dipole moment falls exponentially in N^4 , with N the number of axions. In particular, this contribution is negligible in all models in our ensemble with $N > 17$. We interpret this result as a consequence of large N effects in the geometry that create hierarchies in instanton actions and also suppress the ultraviolet cutoff. We also compute the CP breaking due to high-energy instantons in QCD. In the absence of vectorlike pairs, we find contributions to the neutron electric dipole moment that are not excluded, but that could be accessible to future experiments if the scale of supersymmetry breaking is sufficiently low. The existence of vectorlike pairs can lead to a larger dipole moment. Finally, we show that a significant fraction of models are allowed by standard cosmological and astrophysical constraints.

I. INTRODUCTION

Experiments reveal that QCD preserves CP symmetry to very high precision. The action for QCD can in principle include the CP-violating term

$$\mathcal{L} \supset \frac{\theta}{32\pi^2} \epsilon^{\mu\nu\rho\sigma} \text{tr} G_{\mu\nu} G_{\rho\sigma}, \quad (1)$$

in conventions where the quark mass matrix is real. CP is preserved by QCD if $\theta = 0$, and experiments show that $\theta \lesssim 10^{-10}$ in our universe [1–3]. The strong CP problem is the question of why this number is so small.

There are several possible solutions to this problem, but arguably the most viable proposal, and the one that we will consider in this work, is the Peccei-Quinn (PQ) mechanism [4]. In this scenario, the QCD θ -angle is promoted to a dynamical pseudoscalar $\theta(x)$, the *QCD axion*, with a global shift symmetry $\theta(x) \mapsto \theta(x) + \text{const}$ that is broken by QCD instantons.

The potential induced by low-energy QCD instantons is (see e.g. [5])

$$V_{\text{QCD}}(\theta) = \frac{1}{2} \Lambda_{\text{QCD}}^4 \frac{m_u m_d}{(m_u + m_d)^2} \theta^2 + \mathcal{O}(\theta^4), \quad (2)$$

where $\Lambda_{\text{QCD}}^4 = f_\pi^2 m_\pi^2$, $f_\pi \approx 92$ MeV and $m_\pi = 135$ MeV are the pion decay constant and mass, and m_u and m_d are the u and d quark masses. Provided that (2) is the only significant potential term for θ , minimizing the potential gives $\langle \theta \rangle = 0$, and the strong CP problem is solved: the physical QCD θ -angle θ is zero, in accord with experiments.

The *Peccei-Quinn quality problem* is the fact that sufficiently large breaking of the QCD axion shift symmetry

in another sector can perturb the minimum away from $\theta = 0$. Without loss of generality, the potential for θ given as a linear combination of N pseudoscalar fields ξ^A , $A = 1, \dots, N$ can be written as

$$V(\xi^A) = V_{\text{QCD}}(\theta) + V_{\text{hidden}}(\xi^A), \quad (3)$$

with $V_{\text{QCD}}(\theta)$ given by (2). If the hidden sector potential $V_{\text{hidden}}(\xi^A)$ is such that the minimum of $V(\xi^A)$ occurs at $|\langle \theta \rangle| \gtrsim 10^{-10}$, we say that the PQ solution of the strong CP problem is *spoiled*. The goal of this work is to assess the PQ mechanism in string theory by calculating V_{hidden} in a large class of string compactifications.

The first part of this endeavor concerns the ultraviolet (UV) completion of QCD. We note that θ can receive a contribution from QCD instantons whose zero modes are lifted by higher-dimension fermionic operators, for example four-quark operators. We compute these effects, making three assumptions: the Standard Model is completed into a supersymmetric theory at some high scale $M_{\text{SUSY}} > 1$ TeV; the charged spectrum of QCD is given by Standard Model quarks and their superpartners up to a sufficiently high scale (though we will also comment on the inclusion of vectorlike pairs); and CP is not an approximate symmetry of the ultraviolet theory. We find that although high-scale QCD instantons do not spoil the Peccei-Quinn solution, for sufficiently low supersymmetry breaking their contribution to θ is larger than that of CP-violating effects in the weak sector of the Standard Model. In future experiments on the neutron electric dipole moment, such effects might be detectable.

In the second part of this work, we turn to effects of the UV completion in quantum gravity. Preserving the success of the PQ mechanism amounts to requiring that the breaking of the axion shift symmetry by the hidden sector

potential $V_{\text{hidden}}(\xi^A)$ is at least 10^{-10} times smaller than the breaking by low-energy QCD. This means that simply imposing a weakly broken global shift symmetry onto the effective field theory cannot by itself solve the strong CP problem. In particular, even Planck-suppressed operators can spoil the PQ mechanism, and so a full resolution of the problem requires information about quantum gravity contributions to axion potentials [6–9]. In many corners of string theory, axions — i.e. pseudoscalar fields with approximate shift symmetries — are ubiquitous, and in many constructions the existence of an axion that couples to QCD is automatic. Moreover, in vacua of string theory that furnish weakly-coupled UV completions of gravity, such as weakly-curved compactifications with small values of the string coupling, one can — and we will — compute the leading quantum gravity contributions to the axion potential. String theory thus provides a natural framework for assessing the quality problem.

We study the Peccei-Quinn quality problem in a well-understood class of solutions of string theory: compactifications of type IIB string theory on O3/O7 orientifolds of Calabi-Yau threefold hypersurfaces in toric varieties. In this setting, we suppose that QCD is realized on a stack of D7-branes wrapping a four-cycle D_{QCD} , so that the QCD axion θ arises from reduction of the Ramond-Ramond four-form C_4 on D_{QCD} . The issue is that θ obtains a potential not only from QCD instanton effects, but also from *stringy instantons*: that is, Euclidean D3-branes (or D7-brane gaugino condensates) wrapping various other four-cycles in the manifold. Thus, the quality problem manifests itself through *D-brane instantons that generate the hidden-sector potential for the axions ξ^A* , $A = 1, \dots, N$. Solving the quality problem in this context then amounts to ensuring that these instanton effects have negligible impact on the minimum of the QCD axion potential.

The sizes of the instanton couplings in a given compactification can be obtained, up to a few caveats that we will discuss, by computing the volumes of four-cycles in the internal space: in particular, instantons wrapping large cycles give small couplings. In this work, by computing the volumes of four-cycles in an ensemble of more than 100,000 compactifications, we obtain upper bounds on the CP-breaking couplings encoded in $V_{\text{hidden}}(\xi^A)$. With this in hand, we compute the value of θ for each model.¹

We find that the D-brane instanton contribution to $\langle\theta\rangle$, and hence to the neutron electric dipole moment, falls approximately as $\exp(-cN^4)$ with $c \approx 1.8$. In 99.7% of all the models in our ensemble, and in particular in every

model with $N > 17$, *D-brane instantons do not spoil the PQ mechanism*.

Whether the high-energy QCD instantons discussed above can themselves spoil the PQ mechanism, even when D-brane instantons do not, depends on the spectrum of matter charged under QCD. Counting vectorlike pairs is an often-subtle problem (see e.g. [14]) that is beyond the scope of this work, so we make conditional statements in terms of the number n of such pairs. For $n = 0$, high-energy QCD instantons are harmless, and the PQ mechanism solves the strong CP problem whenever D-brane instanton contributions can be neglected — which is almost always the case in our ensemble. For $n > 0$ vectorlike pairs of mass M_V , high-energy QCD instantons make larger contributions to the neutron electric dipole moment, and the strong CP problem is solved for some but not all values of n and M_V .

The same geometric data needed to characterize the PQ mechanism also provide information about the relic abundances of the various species of axion dark matter, and about the couplings of light axions to photons. Using a simplified model for the dark matter relic abundances and axion-photon couplings, we analyze the astrophysical and cosmological constraints on the axions in our ensemble. We find that for theories in this landscape with $26 \leq N \leq 433$, at least 50% of models are allowed both by astrophysical constraints and by bounds on the axion dark matter relic abundance.

The structure of this paper is as follows. In §II, we explain the Peccei-Quinn quality problem in general terms. In §III, we compute the contribution of high-scale QCD instantons to the QCD axion potential in field theory. We show that these effects do not spoil the Peccei-Quinn mechanism in the absence of vectorlike pairs. In §IV, we describe the incarnation of the Peccei-Quinn quality problem in type IIB compactifications in which the PQ axion descends from the Ramond-Ramond four-form. We summarize the data that one needs to obtain from string theory in order to determine whether the strong CP problem is solved via the PQ mechanism. We then give more details in the specific context of orientifolds of Calabi-Yau threefold hypersurfaces in toric varieties, and describe how we construct an ensemble of axion effective theories. In §V, we present the results from a large ensemble of compactifications of type IIB string theory on orientifolds of Calabi-Yau threefold hypersurfaces. We argue that the strong CP problem is generically solved in this class of theories. We also identify the range of parameter space in our ensemble that is allowed by cosmological and astrophysical bounds. We conclude in §VI. In Appendix A, we detail a fast algorithm to compute masses and decay constants in our models. In Appendix B, we describe two classes of contributions to the scalar potential that we neglect in the main analysis, and we justify the omission of these terms.

¹ The data of a model includes specification of a Calabi-Yau compactification at a given point in its moduli space, but our results turn out to be rather insensitive to the choice of point (see §IV.2), and thus independent of the interplay with moduli stabilization discussed e.g. in [10–13].

II. THE PECCEI-QUINN QUALITY PROBLEM

We consider a set of N axions ξ^A , $A = 1, \dots, N$, with the QCD axion a particular linear combination thereof, $\theta := q_A \xi^A$. The scalar potential can generally be written as

$$V(\theta, \xi^A) = V_{\text{QCD}}(\theta) + V_{\text{hidden}}(\xi^A), \quad (4)$$

where V_{hidden} encodes all contributions to the scalar potential *except* the infrared (IR) contribution from QCD encoded in $V_{\text{QCD}}(\theta)$.

One can think of V_{hidden} as a set of operators that break the continuous shift symmetries of the axions. To illustrate this, we temporarily suppose, as originally envisioned in [4], that the axions ξ^A are the $U(1)$ phases of complex fields Φ^A that obtain non-zero vevs $\Phi^A = f^A e^{i\xi^A}$ from a symmetry breaking potential $V_{\text{symm}}(\Phi^A, \bar{\Phi}^A) \equiv V_{\text{symm}}(|\Phi^A|^2)$ that is independent of the axions. The axions receive a potential e.g. from higher-dimension shift-symmetry breaking operators. In the simplest scenario of a single axion θ , its shift symmetry is broken by operators such as

$$V_{\text{hidden}} \supset \sum_n c_n \frac{\Phi^n}{M_{\text{pl}}^{n-4}} + c.c. \quad (5)$$

Here, without further assumptions, the contributions to the scalar potential from higher-dimension operators scale polynomially in f/M_{pl} . If we take $f \sim 10^{10}$ GeV in order to comply with standard cosmological bounds [15–17], then we see that the operators in (5) must be suppressed up to and including $n = 10$ (we will explain more precisely exactly how much suppression is needed in what follows). For larger decay constants, even higher order operators must be suppressed. From this point of view, the strong CP problem gets traded for an (arguably just as bad) Peccei-Quinn quality problem.

In grappling with the PQ quality problem, one ought to contend with two possibilities: first, that the UV-sensitive shift symmetry breaking from *small* (i.e. high-energy) gauge instantons in QCD is too large; and second, that other sources of shift symmetry breaking, unrelated to the QCD gauge instanton, perturb the minimum of the low-energy potential (2).

The issue of higher-order shift symmetry breaking operators is abated, though not eradicated, in string theory [18]. This is because in certain well-studied corners of the string landscape, the axion shift symmetries are exact in perturbation theory, and are broken only by non-perturbative effects. In such theories, coefficients such as c_n in (5) are exponentially suppressed, in terms of instanton actions that are large at weak coupling. One might hope that in some cases, this suppression will suffice to allow the PQ mechanism to survive. To see if this hope is realized, we will compute these coefficients in a large ensemble of string compactifications and explicitly check whether the shift symmetry is in fact preserved to the high degree necessary.

To set the stage, we first consider a four-dimensional effective field theory with a large UV cutoff M , containing in particular N axion degrees of freedom ξ^A , and a UV version of the Standard Model, but with all stringy degrees of freedom integrated out. As such, the UV Lagrangian will contain an explicit scalar potential $V_{\text{hidden}}(\xi^A)$ for the axions ξ^A generated by physics at scales higher than M . As mentioned above, a key feature of the theories considered in this paper is that the axion shift symmetries are only broken by instanton effects, which preserve a discrete shift symmetry $\xi^A \rightarrow \xi^A + 2\pi k^A$, with $\vec{k} \in \mathbb{Z}^N$. An instanton is naturally labeled by a *charge* $\vec{q} \in \mathbb{Z}^N$, and by a Euclidean action $S_E^{\vec{q}} > 0$, and in terms of these the scalar potential generated by UV instantons is written as

$$\begin{aligned} V_{\text{hidden}}(\xi^A) &= M^4 \sum_{\vec{q} \in \mathbb{Z}^N} A_{\vec{q}} e^{-S_E^{\vec{q}}} \left(1 - \cos(\vec{q} \cdot \vec{\xi} - \varphi_{\vec{q}})\right) \\ &\equiv \sum_{\vec{q} \in \mathbb{Z}^N} \Lambda_{\vec{q}}^4 \left(1 - \cos(\vec{q} \cdot \vec{\xi} - \varphi_{\vec{q}})\right), \end{aligned} \quad (6)$$

with one-loop determinants $A_{\vec{q}} > 0$, and UV CP-breaking phases $\varphi_{\vec{q}}$. Thus we see that any discussion of the PQ quality problem must contend with the possibility that the QCD axion couples to hidden sector axions in a way that shifts the minimum outside of the acceptable range. Without a theory that relates the phases $\varphi_{\vec{q}}$ in such a way that all of the terms in (6) are harmless, these contributions of UV instantons pose a serious threat to the viability of the PQ mechanism.²

Of particular importance in (6) is the charge of the QCD instanton, denoted \vec{q}_Q , which defines the QCD θ -angle via the relation $\theta_Q := \vec{q}_Q \cdot \vec{\xi}$. Crucially, in order for the PQ mechanism to solve the strong CP problem, the UV potential $V_{\text{hidden}}(\xi^A)$ must leave at least one linear combination of axions sufficiently light so that the low energy QCD potential is the dominant contribution to the scalar potential along the light direction. In this situation the QCD θ -angle will remain close to zero.

The contribution to the θ -angle from instantons is easily estimated as

$$\Delta\theta \simeq \frac{\Lambda_{\vec{q}_c}^4}{\Lambda_{\text{QCD}}^4}, \quad (7)$$

where \vec{q}_c is the charge of the instanton with smallest action such that the linear span of \vec{q}_c and all instantons more dominant than \vec{q}_c contains \vec{q}_Q . The strong CP problem is solved if $\Delta\theta < 10^{-10}$, i.e. if

$$\frac{M^4}{M_{\text{pl}}^4} A_{\vec{q}_c} e^{-S_E^{\vec{q}_c}} \gtrsim e^{-S_{\text{crit.}}}, \quad S_{\text{crit.}} \approx 200. \quad (8)$$

² A first-principles calculation of the $\varphi_{\vec{q}}$ is beyond the scope of this work, but the phase associated to the QCD-generated potential arises from the phases in the quark masses, and we find it implausible that such a quantity could directly relate to the phases of (all) the high-scale stringy instantons we will be considering.

The instantons that enter (6) are of two distinct types: there are high-energy QCD instantons, for which \vec{q} is an integer multiple of \vec{q}_Q and $S_{\vec{q}_Q} = \frac{8\pi^2}{g^2(M)}$, and there are bona fide stringy instantons. Our setup will allow us to treat both types of instantons on the same footing, and the total θ -angle is the sum of the contributions from these two classes of instantons:

$$\Delta\theta = \Delta\theta_{\text{QCD}} + \Delta\theta_{\text{stringy}}. \quad (9)$$

In the next two sections we will compute these effects in turn.

We remark that in view of (8), there appear to be three a priori independent mechanisms to lower the scale of the potential from hidden sector instantons:

1. Lowering the cutoff $M \ll M_{\text{pl}}$ suppresses the instanton potential, and it has been argued that this effect occurs naturally in gravitational theories with many light degrees of freedom [19–21]. Indeed, we will find very low UV cutoffs in controlled solutions when the number of axions is large (cf. [22]).
2. Raising the instanton actions $S_E^{\vec{q}}$ would suppress the axion potential parametrically. However, scaling up instanton actions homogeneously by an overall dilatation of the compact space is incompatible with keeping the QCD gauge coupling at its observed value.
3. The one-loop factors $A_{\vec{q}}$ are parametrically small if the SUSY breaking scale is low. However, the running of the QCD coupling is also affected by the SUSY breaking scale, and it will turn out that overall the axion potential is more suppressed when the SUSY breaking scale is *high*: see (21).

III. QCD INSTANTONS AND $\Delta\theta$

We begin by studying CP-breaking effects resulting from QCD gauge instantons. At first glance this might seem to be a pointless exercise, because it is well-known that in a vectorlike and CP-invariant theory, such as that defined by the renormalizable QCD Lagrangian, CP is not broken spontaneously [23, 24]. However, already in the (renormalizable) Standard Model, QCD inherits small violations of CP from the weak interactions [25], as we will review momentarily. Furthermore, and more importantly for our purposes, the low energy potential for the QCD axion generally contains contributions from *small* instantons that are sensitive to CP breaking at very high scales $M \gg 1$ TeV, which remain completely unexplored by the LHC.³

³ The high-energy tail of the instanton potential has been analyzed in various scenarios for beyond the Standard Model physics in the literature [26–31].

Throughout most of this paper we will make the following assumptions about physics between the weak scale and the Planck scale:

- (i) The Standard Model is completed into a supersymmetric theory at a scale $M_{\text{SUSY}} > 1$ TeV.
- (ii) Except for superpartners, there are no light states charged under QCD.
- (iii) Above a UV threshold M , CP is no longer an approximate symmetry.

We will refer to (i) and (ii) collectively as the *SUSY desert assumptions*. Assumption (iii) is broadly applicable in string theory, as emphasized particularly in [18]. On general grounds, it is expected that quantum gravity breaks all global symmetries in the UV; and indeed, in string compactifications with fluxes, requiring all Wilson coefficients of higher-dimension operators to be real would amount to an infinite amount of tuning imposed on the finite set of discrete data defining a model. Finally, we note that if (iii) were not to hold, there would be no strong CP problem in the first place.

At various stages we will relax assumption (ii) to

- (ii*) QCD remains asymptotically free at high energies.

This assumption allows up to *three* vectorlike pairs in the charged spectrum, in addition to the superpartners of the Standard Model quarks.

With these assumptions in hand, we can examine the instanton-generated axion potential. First, let us set our notation. We will work in conventions where all Standard Model fermions are left-handed Weyl fermions. We denote by $Q^i \equiv \begin{pmatrix} u_L^i \\ d_L^i \end{pmatrix}$, $i = 1, 2, 3$, the three generations of quarks in the $(3, 2)_{\frac{1}{6}}$ of the Standard Model gauge group $SU(3) \times SU(2)_L \times U(1)_Y$, by $U^i \equiv (u_R^i)^c$ the up-type $SU(2)$ singlets in $(\bar{3}, 1)_{-\frac{2}{3}}$, by $D^i \equiv (d_R^i)^c$ the down-type $SU(2)$ singlets in $(\bar{3}, 1)_{\frac{1}{3}}$, and by H the Higgs scalar in $(1, 2)_{\frac{1}{2}}$. For each quark generation we define an axial $U(1)_A^i$ such that all quarks of the given generation have charge one, and all others are singlets. The axial $U(1)$ s are anomalous, i.e. under

$$U(1)_A^i : (Q^i, D^i, U^i) \longrightarrow e^{i\alpha_i}(Q^i, D^i, U^i), \quad (10)$$

we have $\theta \rightarrow \theta - 2 \sum_i \alpha_i$. The Yukawa couplings,

$$\mathcal{L}_{\text{int.}} \supset y_{ij}^d D^i H^\dagger Q^j + y_{ij}^u U^i H Q^j + c.c., \quad (11)$$

break the axial $U(1)$ s perturbatively, but below the mass of the Higgs we can just interpret them as Dirac mass terms $m_{ij}^{u,d} = y_{ij}^{u,d} v$ with v the Higgs vev, assign them spurious $U(1)$ charges, and treat the $U(1)$ s as spurious symmetries.

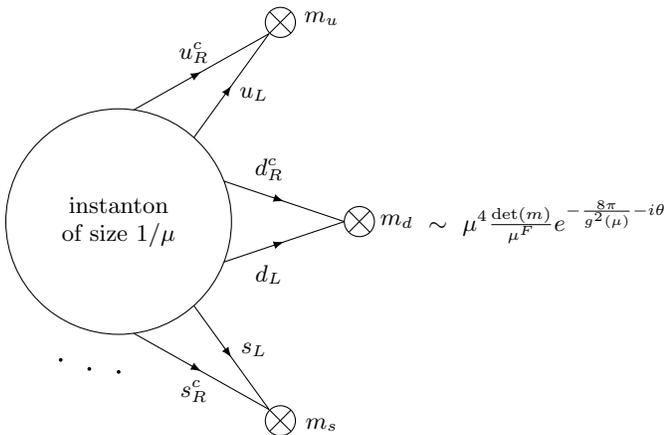


FIG. 1. Lifting of fermion zero modes by mass insertions.

III.1. Infrared contributions

The most straightforward and well-known contribution to the instanton-generated axion potential can be written schematically as an integral

$$\delta V_{\text{inst.}}^m \sim \int \frac{d\mu}{\mu} \mu^4 \frac{\prod_{i=1}^F m_i}{\mu^F} e^{-\frac{8\pi^2}{g^2(\mu)} - i\theta} + c.c., \quad (12)$$

where μ is the inverse size of the instanton, F is the number of quark flavors with mass below μ , and m_i are the quark masses. The mass determinant $\prod_{i=1}^F m_i$ cancels the axial $U(1)$ charge of $e^{-i\theta}$, and equivalently lifts the fermion zero modes of the instanton background (see Figure 1).

Here, as usual, the contribution from large instantons in units of the QCD scale cannot actually be evaluated using the above expression, because QCD is strongly coupled in the IR. This, however, is not much of a practical issue because the IR contribution to the axion potential is determined by the pion mass and decay constant, which are known experimentally. The QCD-induced axion potential reads (see e.g. [32])

$$\delta V^{QCD} = -m_\pi^2 f_\pi^2 \sqrt{1 - \frac{4m_u m_d}{(m_u + m_d)^2} \sin\left(\frac{\theta}{2}\right)^2}. \quad (13)$$

On general grounds, as renormalizable QCD itself does not break CP, the contribution (13) to the axion potential has its minimum at $\theta = 0$, leading to the PQ solution to the strong CP problem [4, 23, 24].

On the other hand, CP-breaking effects beyond QCD, e.g. in the form of higher-dimension multi-fermion operators with complex Wilson coefficients, can also lift fermion zero modes and thus shift the minimum of the axion potential away from zero. In the Standard Model, CP breaking comes from the complex phase in the CKM matrix, and the leading correction with non-trivial CP-breaking phase to the instanton-generated axion potential is depicted in Figure 2 [27].

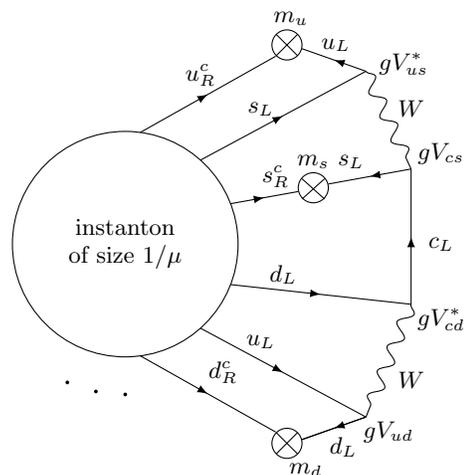


FIG. 2. Lifting of fermion zero modes by CP-breaking weak interactions.

Its effect is still dominated by the IR, where the instanton calculation is not reliable:

$$\delta V_{\text{inst.}}^{\text{CKM}} \sim \int \frac{d\mu}{\mu} \mu^4 \frac{\prod_{i=1}^F m_i}{\mu^F} \times \mu^6 \frac{G_F^2}{m_c^2} V_{ud} V_{cd}^* V_{cs} V_{us}^* e^{-\frac{8\pi^2}{g(\mu)^2} - i\theta} + c.c., \quad (14)$$

which is valid in the range $\mu < m_c$.

However, Standard Model effects induce a dimensionless CP breaking parameter in the chiral effective theory of order [25]

$$\omega := G_F^2 f_\pi^4 |\text{Im}(V_{ud} V_{cd}^* V_{cs} V_{us}^*)| \approx 4 \times 10^{-19}, \quad (15)$$

so one expects $\langle \theta \rangle \sim \omega$ if θ is dynamical, and in the absence of further physics beyond the Standard Model. Thus, if there exists a QCD axion, measuring a QCD θ -angle above the threshold (15) would be a sign of new CP-breaking physics beyond the Standard Model.

III.2. Ultraviolet contributions

Next we estimate the contribution to the axion potential from CP-breaking effects at a high scale M .⁴ In general, one expects that a gauge instanton of size $\sim M^{-1}$ will contribute to the axion potential as

$$\delta V_{\text{inst.}}^{UV}(M) \sim M^4 e^{i\phi_M} e^{-\frac{2\pi}{\alpha_M} - i\theta} + c.c., \quad (16)$$

where $\alpha_M := \frac{g^2(M)}{4\pi}$, and the phase ϕ_M quantifies CP breaking at the scale M . Concretely, CP-breaking Standard Model operators of dimension six and higher (as

⁴ We thank Csaba Csáki and Max Ruhdorfer for useful discussions about this point.

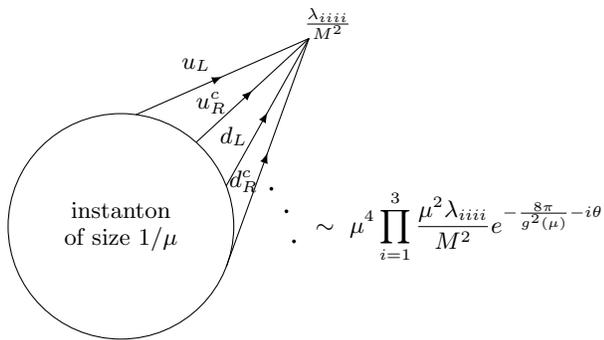


FIG. 3. Lifting of fermion zero modes by CP-breaking dimension six fermion operator.

classified in [33]) can lift all the fermion zero modes. As an example, we consider the dimension six operators

$$\mathcal{O}_6^{ijkl} := \epsilon_{ab} D^i Q^{j,a} U^k Q^{l,b}, \quad (17)$$

where $a, b = 1, 2$ index the fundamental representation of $SU(2)_L$. Insertion of such operators into the effective action as $S_{\text{eff}} \supset \int d^4x \lambda_{ijkl} M^{-2} \mathcal{O}_6^{ijkl} + c.c.$ with non-vanishing dimensionless Wilson coefficients λ_{ijkl} lifts the fermion zero modes of the instanton background: see Figure 3. Importantly, in view of assumption (iii) of §III, the Wilson coefficients λ_{ijkl} will generically have $\mathcal{O}(1)$ complex values.

The resulting potential for the QCD axion takes the schematic form

$$\delta V_{\text{inst.}}^{UV} \sim \int \frac{d\mu}{\mu} \mu^4 \prod_{i=1}^3 \frac{\mu^2 \lambda_{iiii}}{M^2} e^{-\frac{2\pi}{\alpha_\mu} - i\theta} + c.c. \quad (18)$$

We refer to the term (18) as coming from *high-energy QCD instantons*. Unless $\prod_{i=1}^3 \lambda_{iiii} \in \mathbb{R}$, one indeed obtains a term of the form (16) with non-trivial CP breaking phase $\phi_M = \arg(\prod_{i=1}^3 \lambda_{iiii})$.

The magnitude of (18) is controlled by the properties of QCD at high energies. Depending on the matter content of QCD, and assuming $\mathcal{O}(1)$ CP breaking in the UV, small gauge instantons may give a contribution to the axion potential that dominates over the IR potential (13), and thus spoils the PQ solution to the strong CP problem. This problem can be arbitrarily severe if one makes no further assumptions about the QCD spectrum at high energies.

Fortunately, however, this problem of spoliation by high-energy QCD instantons is generally evaded, provided that the charged spectrum is only that of the Standard Model supplemented by superpartners at a SUSY breaking scale M_{SUSY} . In a supersymmetric theory, a gauge instanton has two universal fermion zero modes associated to the breaking of half the supercharges by the instanton. These can only be lifted by a SUSY breaking spurion, i.e. an insertion of the Goldstino mass, which we take to be comparable to the masses of all other superpartners, i.e. of order M_{SUSY} . The instanton contribution

to the scalar potential of (16) is then replaced by⁵

$$\delta V_{\text{inst.}}(M) \sim \frac{8\pi}{\alpha_M} M_{\text{SUSY}} M^3 e^{-\frac{2\pi}{\alpha_M} - i(\theta - \phi_M)} + c.c. \quad (19)$$

Matching the running of the QCD gauge coupling (ignoring threshold corrections) one then finds that⁶

$$\delta V_{\text{inst.}}^{M > M_{\text{SUSY}}} \sim \frac{8\pi}{\alpha_M} \left(\frac{m_Z}{M_{\text{SUSY}}} \right)^3 m_Z^4 e^{-\frac{2\pi}{\alpha_Z} - i(\theta - \phi_M)}, \quad (20)$$

which is independent of the scale M , except for the logarithmic scale dependence of α_M . Here, $m_Z \approx 91$ GeV is the mass of the Z -boson, and $\alpha_Z := \alpha_{m_Z} \approx 0.119$. Thus, upon ignoring an overall $\mathcal{O}(1)$ factor⁷ α_Z/α_M that depends only logarithmically on the value of M due to the running of α between the mass of the Z -boson and the UV scale M , one finds a contribution to the axion potential from high-energy QCD instantons of order $\delta V_{\text{inst.}} \sim 10^{-12} \left(\frac{1 \text{ TeV}}{M_{\text{SUSY}}} \right)^3 \Lambda_{\text{QCD}}^4$. Such a contribution would shift the minimum of the QCD axion potential to

$$\Delta\theta_{\text{QCD}} \sim 10^{-12} \left(\frac{1 \text{ TeV}}{M_{\text{SUSY}}} \right)^3. \quad (21)$$

Thus, given the fact that $M_{\text{SUSY}} \gtrsim 1$ TeV [34], we find that (21) is within the experimentally allowed range.

But importantly, this conclusion is sensitive to our SUSY desert assumptions — (i) and (ii) of §III — that the QCD spectrum is empty up to high scales except for superpartners. To see this, we now relax our assumptions (ii)→(ii*), i.e. we allow for $n = 1, 2$, or 3 vectorlike pairs (in chiral multiplets) with mass M_V in the window $M > M_V > M_{\text{SUSY}}$.⁸ This leads to a softer running of the QCD coupling in the UV, and the right-hand side of (20) gets modified to

$$\frac{8\pi}{\alpha_M} \left(\frac{m_Z}{M_{\text{SUSY}}} \right)^3 \left(\frac{M}{M_V} \right)^n m_Z^4 e^{-\frac{2\pi}{\alpha_Z} - i(\theta - \phi_M)} + c.c. \quad (22)$$

⁵ We have normalized the instanton potential so that it matches the leading contribution to the F-term potential of the four-dimensional supergravity models considered in this paper. In other models the normalization might be slightly different.

⁶ Here, for simplicity we use six ‘active’ quarks down to the Z -boson threshold, ignoring the fact that $m_{\text{top}} > m_Z$: this produces only an $\mathcal{O}(1)$ error.

⁷ In the extreme case where $M \gtrsim M_{\text{GUT}}$ and $M_{\text{SUSY}} = 1$ TeV this means accepting a mild correction factor $\alpha_Z/\alpha_M \approx 4$.

⁸ Although true vectorlike pairs with small masses are arguably non-generic, string theory realizations of the Standard Model often contain extra Green-Schwarz massive $U(1)$ factors [35, 36]. Pairs that are vectorlike with respect to the Standard Model can still be chiral with respect to the extra $U(1)$ factors and have naturally small masses generated by instantons. We thank Jim Halverson for pointing this out to us, and for an insightful discussion about this possibility.

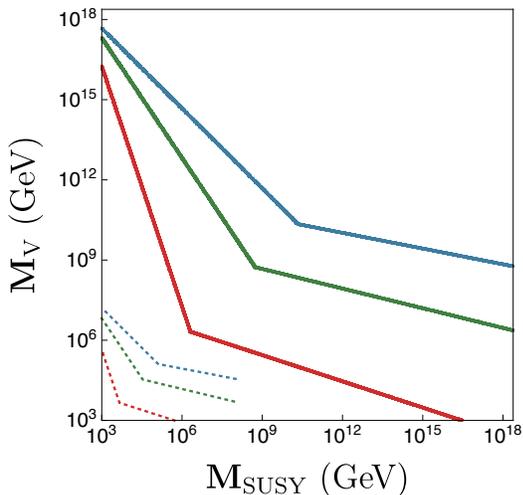


FIG. 4. Lower bound on the vectorlike masses as a function of the scale of supersymmetry breaking, for UV reference scales $M = M_{\text{Pl}}$ (solid) and $M = 10^8$ GeV (dashed). In red is the lower bound for one vectorlike pair, in green for two, and in blue for three. Theories living above these lines have allowed contributions to the θ -angle from high-scale QCD instantons.

For opposite ordering of scales $M > M_{\text{SUSY}} > M_V$, and scalar masses of order M_{SUSY} , one obtains (22) with $M_V \rightarrow M_{\text{SUSY}}^{\frac{1}{3}} M_V^{\frac{2}{3}}$. Again, for large SUSY breaking scale and not too light vectorlike pairs, $\Delta\theta_{\text{QCD}}$ is below the experimental bound, but for low SUSY breaking scale, and at least one light vectorlike pair, the PQ mechanism can be spoiled. We show the allowed region in Fig. 4. Under our assumptions, preserving the PQ mechanism implies that a light vectorlike pair observed at collider experiments would imply an unobservably large SUSY breaking scale, and vice versa.

We therefore find, with hardly any model-dependence, that PQ quality is endangered the least when the SUSY desert assumptions hold. This leads us to a striking conclusion. If the scale of SUSY breaking is relatively low, $M_{\text{SUSY}} \lesssim 100$ TeV, then the CP-breaking effects of high-energy QCD instantons exceed the CP breaking (15) from weak interactions, even if the scale M where CP breaking becomes $\mathcal{O}(1)$ is as high as the Planck scale. Thus, an observable SUSY breaking scale would imply *either* that the QCD θ -angle is greater than the Standard Model expectation computed in (15), *or* that the UV completion preserves CP to a considerable degree. In the rather optimistic scenario that SUSY is ultimately found close to the TeV scale and future experiments constrain the neutron electric dipole moment below the threshold given by (21), one would rule out the entire landscape of solutions considered in this paper. Ultimately, one would either rule out the hypothesis of $\mathcal{O}(1)$ CP breaking in the deep UV, or reintroduce the strong CP problem, without a viable axion solution.

Having understood the contribution to the axion potential from QCD effects at low and high scales, we now

turn to all other contributions to the axion potential that arise in UV-complete models. For this, we will need to invoke concrete UV completions in string theory.

IV. THE PECCEI-QUINN QUALITY PROBLEM IN TYPE IIB STRING THEORY

We will now detail the circumstances in which a model in type IIB string theory possesses a quality problem that obstructs the Peccei-Quinn solution to the strong CP problem. For prior work on this subject, see e.g. [10, 13, 37–40].

IV.1. Type IIB conditions for a PQ quality problem

To begin, we will briefly describe features of low-energy supergravity theories that arise via compactification of type IIB string theory on Calabi-Yau orientifolds. We will focus on expressing the instanton generated potential (6) in terms of geometric data.

In type IIB on a Calabi-Yau orientifold⁹ X (of O3/O7 type) with fluxes, or an F-theory generalization thereof, the effective theory at low energies is well-described by $\mathcal{N} = 1$ supergravity with Standard-Model-like gauge sectors from open string degrees of freedom on seven-branes and their intersection curves (vector multiplets and charged chiral multiplets, respectively), along with N gauge singlet chiral multiplets T_A , $A = 1, \dots, N$, where $N = \dim H^2(X, \mathbb{R})$. The real parts $\tau_A := \text{Re}(T_A)$ are geometric moduli of the threefold X and can be thought of as the volumes of a homology basis of four-cycles $D_A \in H_4(X, \mathbb{Z})$, and the imaginary parts $\xi_A := \text{Im}(T_A)$ are the axion zero modes of the four-form potential C_4 . The (classical) holomorphic Yang-Mills coupling associated with a stack of seven-branes wrapped on a divisor in a class $q^A[D_A]$ is equal to $q^A \tau_A$, and the linear combination $q^A \xi_A$ takes on the role of a dynamical θ -angle. The gauge symmetry $C_4 \rightarrow C_4 + d\Lambda_3$ in ten dimensions implies that the ξ_A receive no scalar potential to any perturbative order in the large volume expansion. Rather, their potential is entirely generated by instantons: specifically, by Euclidean D3-branes wrapped on 4-cycles in the compactification manifold.

The action of a Euclidean D3-brane wrapping an *effective* divisor $\Sigma_{\vec{q}}$ (i.e. a holomorphic four-cycle) in a homology class $[\Sigma_{\vec{q}}] = q^A[D_A]$ is

$$S_E^{\vec{q}} = 2\pi \text{Vol}(\Sigma_{\vec{q}}) = 2\pi q^A \tau_A, \quad (23)$$

⁹ We do not perform an explicit orientifold, but we note that projecting instantons out by orientifolding would only strengthen our results on PQ quality.

and the instanton expansion of the superpotential is, setting $M_{\text{pl}} = 1$,

$$W = W_0 + \sum_{\bar{q}} \mathcal{A}_{\bar{q}} \exp(-2\pi q^A T_A). \quad (24)$$

Here, W_0 is a (generally complex) constant term generated by fluxes, and the sum runs over all effective divisors. Euclidean D3-branes wrapping cycles that are *not* effective break all supercharges, and therefore may at most contribute to the Kähler potential

$$K = K_{\text{tree}}(\tau_A) + K_{\text{pert.}}(\tau_A) + K_{\text{np}}(T_A, \bar{T}_{\bar{A}}). \quad (25)$$

Here, the tree-level Kähler potential is $K_{\text{tree}}(\tau_A) = -2 \log \mathcal{V}$ where \mathcal{V} is the geometric volume of X (as a function of the τ_A), and the perturbative and non-perturbative corrections $K_{\text{pert.}}(\tau_A)$ and $K_{\text{np}}(T_A, \bar{T}_{\bar{A}})$ are not known in general, but are negligible when the curvature of the internal space is small.

For simplicity, let us assume that the τ_A are stabilized, perturbatively, at a point in field space where $\mathcal{A}_{\bar{q}} e^{-2\pi q^A T_A} \ll |W_0| \ll 1$. Then the F-term potential can be written as

$$V_F = V_{\text{pert.}}(\tau_A) + V_{\text{axion}}(\tau_A, \xi_A), \quad (26)$$

with

$$V_{\text{axion}}(\tau_A, \xi_A) = 2m_{\frac{3}{2}} e^{\frac{\kappa}{2}} \text{Re}(\bar{\partial}^A K_{\text{pert.}} \partial_A W_{\text{np}} + \epsilon W_{\text{np}}) + \mathcal{O}\left(W_0^2 e^{-2\pi q^A \tau_A}, e^{-2\pi(q^A + q'^A)\tau_A}\right), \quad (27)$$

where $m_{\frac{3}{2}} := e^{\frac{\kappa}{2}} |W|$ is the gravitino mass, $\epsilon := K^{A\bar{B}} K_A K_{\bar{B}} - 3$, and without loss of generality we have taken $W_0 \in \mathbb{R}$. In Calabi-Yau compactifications the tree-level Kähler potential K_{tree} satisfies the relations $\epsilon = 0$ and $\bar{\partial}^A K_{\text{tree}} = -2\tau^A$, so it is justified to approximate the axion potential (6) using

$$\Lambda_{\bar{q}}^4 \approx 8\pi m_{\frac{3}{2}} \frac{\bar{q} \cdot \bar{\tau}}{\mathcal{V}} |\mathcal{A}_{\bar{q}}| e^{-2\pi \bar{q} \cdot \bar{\tau}}, \quad (28)$$

i.e. one considers the axion potential generated by irreducible BPS instantons. Note that multi-instanton terms are strictly subleading to terms of the form (28)—therefore, a quality problem introduced by instantons of the form (28) cannot be remedied by the inclusion of multi-instanton terms, and conversely if no quality problem is present at the level of (28), one cannot be introduced by these subleading effects. Crucially, for any point in field space where the above approximations hold we can compute the leading instanton corrections to the scalar potential up to the standard $\mathcal{O}(1)$ uncertainty in the holomorphic one-loop Pfaffians $\mathcal{A}_{\bar{q}}$. Even if some or all of the instanton terms instead dominate over a small constant term W_0 , but remain small themselves, on general grounds one may still use the estimate $\Lambda_{\bar{q}}^4 \sim m_{\frac{3}{2}}/\mathcal{V} \times e^{-2\pi \bar{q} \cdot \bar{\tau}}$ except that $m_{\frac{3}{2}}$ itself is dominated

by instanton contributions. Thus, as long as the scales $\Lambda_{\bar{q}}$ in (28) are expressed in terms of the gravitino mass instead of the value of the classical flux superpotential W_0 , we may estimate the axion potential using (28).

On the other hand, if the overall superpotential is large, non-BPS instantons can dominate over BPS instantons, and no general methods now exist for computing the axion potential (see however [41, 42] and Appendix B for comments on this regime). For this reason, we will by necessity take the SUSY breaking scale to be small.

One might suspect that the assumption of a low SUSY breaking scale, made here for reasons of computational control, severely limits what can be said about the prevalence of a PQ quality problem in our framework. However, as we have seen in the previous section, under some mild assumptions about the matter spectrum of QCD, the PQ quality problem tends to get more severe the *lower* the SUSY breaking scale. Thus, we see little reason to expect that the PQ quality problem is present when the SUSY breaking scale is high. In other words, it appears as a fortunate coincidence that one can draw fairly general conclusions about the strong CP problem by studying instanton corrections to the *superpotential* (which is comparatively accessible by direct computation), because the problem becomes milder as one leaves the regime where superpotential terms are dominant.

IV.2. Generating models

In order to assess PQ quality in a *fully explicit* landscape of models of the above type one would have to generate a rather impressive set of data:

1. A large set of Calabi-Yau threefolds.
2. One or more consistent O3/O7 orientifold involutions for each threefold.
3. For each geometric model, choices of seven-branes and fluxes that generate a realistic Standard Model, with somewhat small scale of SUSY breaking, moduli stabilization, the correct dark matter abundance, etc.

For reasons of time we will not pursue achieving this here. Instead, we will content ourselves with a large set of Calabi-Yau threefolds from the Kreuzer-Skarke list [43]; a set of effective divisors that could in principle host QCD; and a point in the geometric moduli space of each Calabi-Yau where the instanton expansion of the superpotential converges and the QCD gauge coupling comes out right. We do not know which, if any, of these toy models can actually be completed to fully realistic models of particle physics and cosmology. Even so, it would require a quite remarkable coincidence for the divisors hosting QCD, and the points in moduli space where moduli stabilization *actually* occurs to be so special as to invalidate conclusions drawn from the statistics of a large random sample.

A few words are in order about the construction of our ensemble. Our aim is to generate a random sample of Calabi-Yau manifolds, and determine an upper bound on the QCD θ -angle in these geometries. We will realize these manifolds as hypersurfaces in toric varieties, following [44]. The starting point is a four-dimensional reflexive polytope Δ° from the Kreuzer-Skarke list and its set of fine regular star triangulations $\{\mathcal{T}\}$, each defining a generically singular toric fourfold $V_{\mathcal{T}}$ whose generic anti-canonical hypersurface defines a smooth Calabi-Yau threefold $X_{\mathcal{T}}$.¹⁰

Let $\mathcal{K}_{\mathcal{T}}$ be the sub-cone of the Kähler cone of $X_{\mathcal{T}}$ that is inherited via restriction from the Kähler cone of $V_{\mathcal{T}}$, and denote by $\mathcal{K} := \cup_{\mathcal{T}} \mathcal{K}_{\mathcal{T}}$ the *inherited extended Kähler cone*. Intuitively, \mathcal{K} can be thought of as the part of the Calabi-Yau threefold's Kähler moduli space that is describable via embedding into the set of birationally equivalent toric fourfolds.

We obtain a random sample of triangulations using Algorithm #3 of [45]. An implementation of this algorithm is included in `CYTools` [46], via the function `random_triangulations_fair`: we used the parameters `initial_walk_steps=500`, `n_walk=50` and `n_flip=50`.

This method picks out a random Calabi-Yau compactification. Equipped with an explicit compactification manifold, one might expect that all our conclusions will heavily depend on the precise choice of point in its Kähler moduli space. For each direction in the Kähler cone, the overall dilation is fixed by demanding that QCD has the right gauge coupling (to be described in detail momentarily), so our results could, a priori, be sensitive to the choice of ray in $\mathcal{K}_{\mathcal{T}}$. However, at large $h^{1,1}$, the inherited extended Kähler cone contains exponentially many phases $\mathcal{K}_{\mathcal{T}}$ connected to each other via bi-stellar flips in the triangulation. As a consequence, within each individual phase $\mathcal{K}_{\mathcal{T}}$, divisor volumes are expected to depend only very weakly on the choice of direction in the Kähler cone.

Thus, for simplicity, in each phase we select as reference point a homogeneous dilation of the *tip of the stretched Kähler cone*, t_{tip} , which has been used in various recent studies of the string landscape [22, 45, 47–49], and is defined as the point closest to the origin where all effective curves have volumes ≥ 1 . As explained above, we do not expect this choice to actually matter, but in order to test our expectation quantitatively we have run auxiliary scans using other randomly chosen rays in $\mathcal{K}_{\mathcal{T}}$, obtained via a slight variation of Algorithm #2 of [45], with `n_walk=10`. We indeed find that our results remain the same to good approximation.

Having chosen a one-dimensional ray for each geometric model, we now impose two important constraints:

- The gauge coupling of QCD at this point must match the gauge coupling of QCD in the IR.
- The instanton expansion must be under control.

In order to address the first point, we note that along *any* ray in $\mathcal{K}_{\mathcal{T}}$ and any choice of divisor hosting QCD there exists a point at which it would have the correct volume to reproduce the QCD gauge coupling at low energies. Thus, we need to (a) choose a divisor hosting QCD and (b) homogeneously scale t_{tip} in order to ensure that the high-scale QCD coupling indeed runs to the observed QCD coupling at low energies. For step (a) the toric fourfold provides natural candidates: the intersections of the $h^{1,1}(V) + 4$ toric divisors of V with the Calabi-Yau hypersurface. We will generate a model for a sample of up to five of these. We emphasize that step (b) is sensitive to the running of the QCD coupling in the UV, which, under assumptions (i) and (ii*), is set by the SUSY breaking scale, the number n of vectorlike pairs charged under QCD, and their mass M_V .

The correct homogeneous scaling of the Kähler parameters can be determined from the running of the QCD coupling from the UV matching scale M to a low energy scale, and must be chosen such that the following equation holds,

$$e^{-2\pi \text{Vol}(D_Q)} = \left(\frac{m_Z}{M_{\text{SUSY}}} \right)^3 \left(\frac{M}{\tilde{M}_V} \right)^n \times \frac{m_Z^4}{M^3 M_{\text{SUSY}}} e^{-\frac{8\pi^2}{g^2(m_Z)}}, \quad (29)$$

where D_Q is the divisor hosting QCD, and $\tilde{M}_V := M_V$ for $M_V > M_{\text{SUSY}}$ and $\tilde{M}_V := M_V^{\frac{2}{3}} M_{\text{SUSY}}^{\frac{1}{3}}$ otherwise. For a given ray in the Kähler cone, this is a transcendental equation in one real parameter, which is straightforward to solve numerically. In what follows we will primarily focus on the case $n = 0$ and comment on the generalization to $n > 0$ at the end of §V. Importantly, as PQ quality is endangered most for small SUSY breaking scale, we will set the masses of superpartners to be of order 1 TeV $\approx 10^{-15} M_{\text{pl}}$, roughly as light as allowed with current exclusion limits from the LHC.

After performing the homogeneous rescaling to impose (29), we compute the volumes of all prime toric divisors (cf. Appendix B) in the Calabi-Yau. In order to impose our second constraint, i.e. convergence of the instanton expansion, we will not attempt to assess PQ quality in models with small divisors: such theories are not under computational control, and statements about the θ -angle based on the first few terms in a non-convergent series expansion should not be trusted. Thus, once we have dilated the overall volume to match the QCD gauge coupling at low energies, we will only compute the θ -angle if all divisor volumes are greater than one:

$$\text{Vol}(D_I) \geq 1 \quad \forall I. \quad (30)$$

¹⁰ In this work, we will restrict our ensemble to favorable hypersurfaces for ease of computation. We are not aware of any properties of non-favorable hypersurfaces that would modify our results.

It would clearly be of great interest to repeat the analysis in the regime of small divisor volumes, but such a computation is beyond the scope of this work. One could speculate that such an investigation would reveal that Peccei-Quinn quality is positively correlated with control of the α' expansion, i.e. that the strong CP problem is solved in (most of) the geometric regime, and *not* solved in large portions of the *non*-geometric regime. Such a finding would be a rare example of more desirable phenomenology occurring in a regime that is convenient for theorists. Indeed, this would be an example of the proverbial keys actually being under the lamppost.

Finally, we need to define what precisely we mean by the UV scale M . First, we note that in general it does not correspond to a physical mass threshold but is rather a generally unphysical matching scale chosen such that extrapolation of the four-dimensional effective theory to that scale gives $\frac{4\pi}{g^2(M)} = \text{Vol}(D_Q)$. It is straightforward to estimate this scale by comparing to, for instance, the case of gaugino condensation in pure super-Yang-Mills obtained from \mathfrak{c} seven-branes on a rigid divisor D . In the supergravity theory one expects a superpotential $W_{np} \supset \mathfrak{c}\mathcal{A}_D e^{-2\pi T_D}$ where T_D is the complexified volume of the cycle, and \mathcal{A}_D is a one-loop Pfaffian that depends only on complex structure moduli and the dilaton, and thus generically takes $\mathcal{O}(1)$ values. The well-known gaugino condensation superpotential in four-dimensional super Yang-Mills is $\mathfrak{c}\Lambda^3$ [50, 51] where $\Lambda^3 = M^3 e^{-2\pi T/c}$ and this gets identified with $e^{\frac{K}{2}} M_{\text{pl}}^3 \mathfrak{c}\mathcal{A}_D e^{-2\pi T_D}$ where $K = -2\log(\mathcal{V}) - \log(1/g_s)$ is the Kähler potential. We thus find

$$M \simeq e^{\frac{K}{6}} M_{\text{pl}} = \frac{g_s^{\frac{1}{6}}}{\mathcal{V}^{\frac{1}{3}}} M_{\text{pl}} = M_s^{\frac{2}{3}} M_{\text{pl}}^{\frac{1}{3}}, \quad (31)$$

where M_s is the string scale. Note that $M_{KK} < M_s < M < M_{\text{pl}}$, so the matching scale M is actually larger than the naïve cutoff of the four-dimensional effective theory, namely the Kaluza-Klein scale. For this work, we will assume only moderately small string coupling, so we may choose $M \simeq M_{\text{pl}}/\mathcal{V}^{\frac{1}{3}}$.

V. RESULTS

We generated an ensemble of 32,040 manifolds with Hodge numbers ranging from $h^{1,1} = 2$ to $h^{1,1} = 491$, with up to 10 polytopes per $h^{1,1}$ and up to 10 triangulations per polytope, along with the values of the Kähler moduli at the tip of the stretched Kähler cone in each geometry. For each such point, we considered hosting QCD on up to 5 prime toric divisors.¹¹ As stated in §IV.2, we chose the scale of supersymmetry breaking to be 1 TeV in order to

put the Peccei-Quinn mechanism in the most danger. In total, our ensemble consists of 136,659 distinct models.

In Fig. 5, $\Delta\theta_{\text{stringy}}$ and $\Delta\theta_{\text{QCD}}$ are plotted as functions of the number of axions $N \equiv h^{1,1}$, in blue and red, respectively. Each point represents one model. Recall that the total upper bound on the θ -angle is

$$\Delta\theta = \Delta\theta_{\text{QCD}} + \Delta\theta_{\text{stringy}}, \quad (32)$$

as in (9).

We draw the reader's attention to two regions on this plot. The region shaded red is excluded by experimental bounds on the neutron electric dipole moment. The fact that θ is found not to be in this region is the strong CP problem. In the region shaded blue, the θ -angle is smaller than the floor (15) set by CP-breaking in the weak sector [25]. In the narrow band between the red and blue regions, one could hope to measure a θ -angle that is not currently excluded, but is governed by physics beyond the Standard Model.

In a small fraction of the models with $N \lesssim 20$, $\Delta\theta_{\text{stringy}}$ is large enough to spoil the PQ mechanism: specifically, of the 6,716 models in the ensemble with $N \leq 17$, 414 (8%) have $\Delta\theta_{\text{stringy}} > 10^{-10}$. However, $\Delta\theta_{\text{stringy}}$ quickly becomes utterly negligible as N increases: we find

$$\Delta\theta_{\text{stringy}} \propto e^{-cN^4}, \quad (33)$$

with $c \approx 1.8$.

We therefore conclude

$$\Delta\theta_{\text{stringy}} \ll \Delta\theta_{\text{QCD}} < 10^{-10}. \quad (34)$$

The first inequality in (34) holds for $N \gg 1$, but is independent of the number n of vectorlike pairs. The second inequality is independent of N but is modified if $n > 0$ (see §III.2).

Fig. 5 shows that for $N > 17$, $\Delta\theta_{\text{QCD}}$ is the dominant contribution to the θ -angle.¹² Notably, at $M_{\text{SUSY}} = 1$ TeV, $\Delta\theta_{\text{QCD}}$ is larger than the floor (15) from CP-breaking by the weak interactions. Even so, as explained in §III.2, *in the absence of vectorlike pairs* the contribution $\Delta\theta_{\text{QCD}}$ is always below the experimental bound on θ , and the only effect that can shift the minimum of the QCD axion potential far enough to spoil the PQ mechanism is $\Delta\theta_{\text{stringy}}$.

We conclude that theories in this landscape with $N \gtrsim 20$ axions and no vectorlike pairs are exponentially likely to realize the Peccei-Quinn solution of the strong CP problem. This is one of our main results.

are not enough divisors whose volumes can generate the correct QCD gauge coupling via an overall dilation without shrinking any other divisors to have volumes less than one.

¹¹ For some values of $h^{1,1}$ there are fewer than 10 polytopes, and for some polytopes there are fewer than 10 triangulations. The number of models per Calabi-Yau can be less than five if there

¹² For small values of N , it is sometimes the case that $\Lambda_{\text{QCD}}^4 \gg \Lambda_{\text{QCD}}$. In these cases, we select a random value for $\Delta\theta_{\text{stringy}}$ between 0 and 2π .

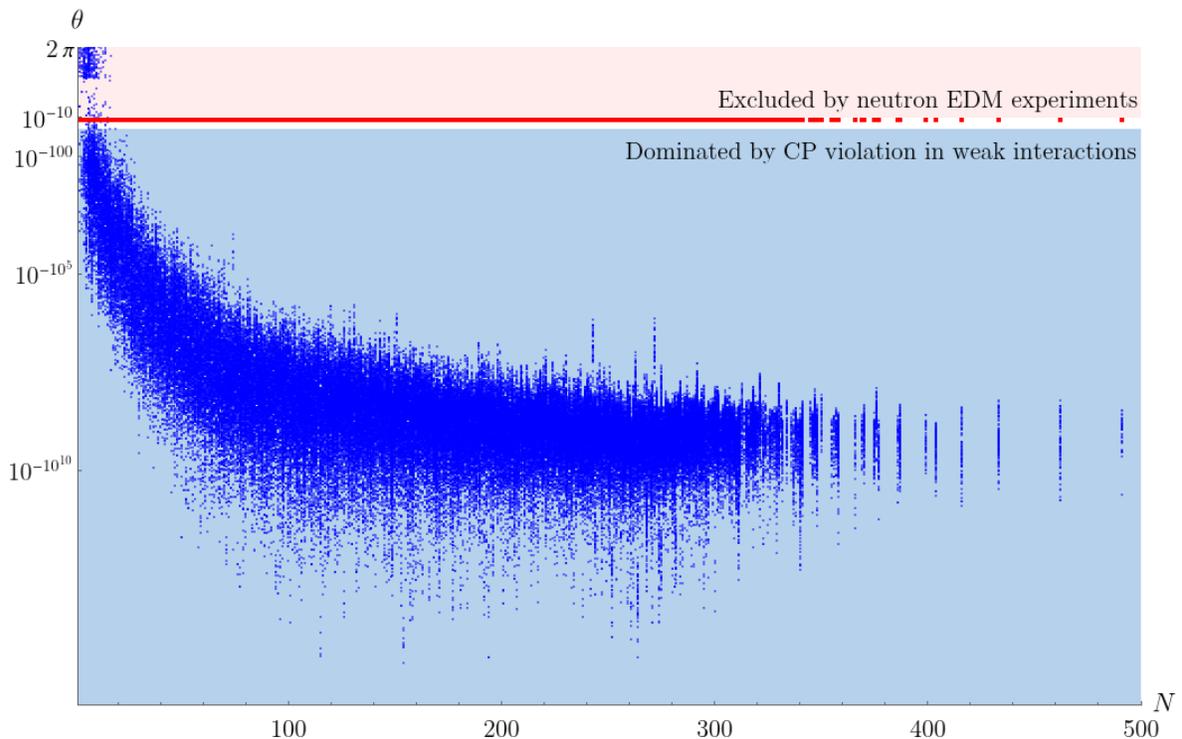


FIG. 5. Contributions to the θ -angle as functions of the number of axions, N . For each model — a specific Calabi-Yau with a choice of a divisor hosting QCD — there is a red point showing $\Delta\theta_{\text{QCD}}$ and a blue point showing $\Delta\theta_{\text{stringy}}$, cf. (9). 136,659 models are shown. The pink region indicates the region $\theta > 10^{-10}$, which is excluded by bounds on the neutron electric dipole moment. The blue region shows the parameter space $\theta < 10^{-19}$, which is below the floor set by CP-breaking effects in the weak sector of the Standard Model.

We end with two remarks. First, we comment on the quality problem in the case that the non-perturbative effects that contribute to (24) are not Euclidean D3-branes, but rather strong gauge dynamics on D7-branes stacked on various cycles. Gaugino condensation on a four-cycle with volume T wrapped by \mathfrak{c} D7-branes would produce a contribution $\propto e^{-2\pi T/\mathfrak{c}}$ to (24), thereby enhancing the scale of V_{hidden} . If \mathfrak{c} could be arbitrarily large, then this enhancement could be great enough to spoil the PQ mechanism. However, in a typical model, one might expect to find $\mathfrak{c} \sim 6$ — see e.g. [52, 53]. For any comparably small value of \mathfrak{c} , the conclusions we obtained about the Peccei-Quinn mechanism remain unchanged: the values of $\Delta\theta_{\text{stringy}}$ that we found are so small (see Fig. 5) that the qualitative conclusions are robust to $\mathcal{O}(1)$ changes in divisor volumes. Even so, the precise value of N for which good quality is guaranteed may shift. Furthermore, if \mathfrak{c} were large, it would appear reasonable to require that all divisor volumes are $\gtrsim \mathfrak{c}$ instead of imposing (30), effectively removing the parameter \mathfrak{c} from the problem.

Second, if there are vectorlike pairs charged under QCD, some portions of the space parameterized by the SUSY breaking scale and the masses of vectorlike pairs become excluded, as explained in §III.2. Even so, *provided* these parameters take values in the allowed regime shown in Figure 4, our results obtained under the as-

sumption of no vectorlike pairs immediately carry over: for each model a change in the above parameters simply translates into a dilation of the overall volume via the matching condition (29), and the relative cycle sizes are invariant. Therefore, if stringy instantons are subleading to the high scale QCD instantons in a model with no vectorlike pairs, they will remain so in a corresponding model with $n > 0$ pairs.

V.1. Cosmological and astrophysical bounds

Up to this point we have analyzed the strong CP problem in an ensemble of axion theories, without imposing experimental constraints other than the upper bound on the neutron electric dipole moment and the lower bound on the scale of supersymmetry breaking. We now consider constraints on this ensemble from the measured dark matter relic density, and from stellar emission of axions and axion-like particles (ALPs). Such limits are not the main purpose of this work, and so our treatment will be schematic, and follow well-known lines. A comprehensive analysis would be a useful target for future work.

We begin by discussing the physics considerations that lead to a bound on the axion dark matter relic density.

We assume that the initial values of the dimensionless axion fields are of order unity, and we consider each species that is stable on cosmological timescales, i.e. with (cf. e.g. [54])

$$\frac{\tau_{\text{ALP}}}{\tau_{\text{universe}}} = 10^{-15} \left(\frac{\text{GeV}}{m_a} \right)^3 \left(\frac{f_a}{10^{12} \text{ GeV}} \right)^2 \gtrsim 1, \quad (35)$$

where τ_{universe} is the age of the universe, m_a is the ALP mass, and f_a is the ALP decay constant. The fractional energy density of one¹³ such ALP is [57]:

$$\Omega_a \approx \frac{1}{6} (9\Omega_r)^{3/4} \left(\frac{m_a}{H_0} \right)^{1/2} \left(\frac{\phi_i}{M_{\text{pl}}} \right)^2, \quad (36)$$

where Ω_r is the fractional energy density due to radiation, and ϕ_i is the initial axion field displacement. The observed relic energy density of dark matter $\Omega_{DM} h^2 = 0.12$ then leads to a bound on the ALP mass:

$$m_a \lesssim \frac{10^{37} \text{ GeV}^5}{f_a^4}. \quad (37)$$

We will consider a model to be cosmologically viable if all the axions in the theory that satisfy (35) also satisfy (37).

For the specific case of the QCD axion, one can use $m_a \sim \frac{m_\pi f_\pi}{f_a}$ in (37) to find the standard bound

$$f_a \lesssim 10^{12} \text{ GeV} \quad (38)$$

on the QCD axion decay constant, for an initial misalignment angle of order unity.

Observational constraints on the axion-photon coupling, $g_{a\gamma\gamma}$, provide lower bounds on the QCD axion decay constant. One of the most stringent bounds comes from the CERN Axion Solar Telescope (CAST) [58], which finds

$$g_{a\gamma\gamma} \lesssim 0.66 \times 10^{-10} \text{ GeV}^{-1} \quad (39)$$

for axions with masses less than 0.02 eV.¹⁴

We estimate the coupling $g_{a\gamma\gamma}$ by taking the axion that couples to $F_{\mu\nu} F^{\mu\nu}$ to be a general linear combination of all mass and kinetic eigenstate axions in the theory, with $\mathcal{O}(1)$ coefficients. Then we define

$$g_{a\gamma\gamma, \text{ eff}} = \frac{\alpha_{\text{EM}}}{2\pi} \sqrt{\sum_{i|m_i < 0.02 \text{ eV}} \frac{1}{f_i^2}} \quad (40)$$

¹³ We do not incorporate mixings of axion species, though it would be interesting to do so. For studies of effects of axion mixing on dark matter abundance and on axion-photon conversion, see [55] and [56], respectively.

¹⁴ Other axion searches are sensitive to lower values of the axion-photon coupling in other regions of parameter space [59–64], but we leave a detailed analysis of these bounds to future work.

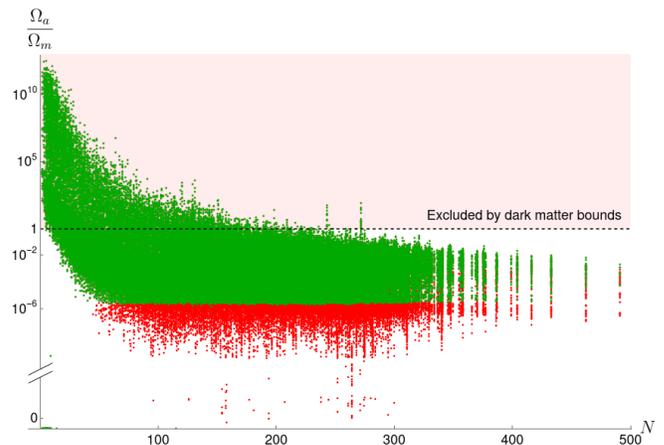


FIG. 6. The fractional abundance of axion dark matter as a function of N . The dotted line at 1 indicates the place where axion dark matter constitutes the entirety of the observed dark matter abundance. Values higher than this are excluded. Green points are allowed by the constraints from CAST, while red points (17%) are not.

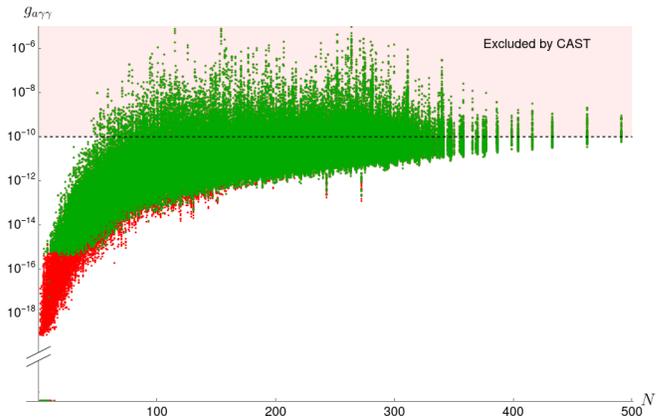


FIG. 7. The effective axion-photon coupling as a function of N . Points above the line are excluded by CAST [58]. Green points are allowed by dark matter relic abundance constraints, while red points (9.4%) are not.

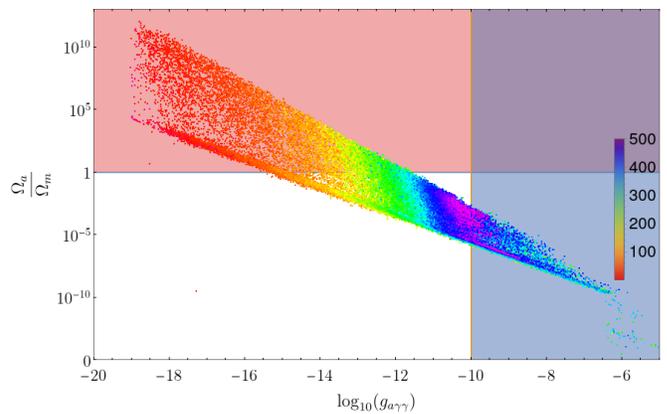


FIG. 8. Dark matter relic density versus effective $g_{a\gamma\gamma}$, color-coded by N . Red points correspond to low N , while purple points correspond to large N .

and compare this with observational bounds.

Cosmological and astrophysical constraints on our ensemble are shown in Figs. 6-8. Fig. 6 shows the fractional abundance of axion dark matter as a function of N . Fig. 7 shows the effective $g_{a\gamma\gamma}$ as a function of N . Finally, Fig. 8 displays these data as dark matter relic density versus effective $g_{a\gamma\gamma}$, color-coded by N .¹⁵

Several remarks about these results are in order. First, one can see that an order-one fraction of all the geometries studied are allowed by observational bounds. The dark matter bounds become less stringent as N increases: this is because the Kaluza-Klein cutoff drops with N , so decay constants decrease as well. Conversely, astrophysical bounds become more stringent as N increases. Even so, the models in our landscape are generally not ruled out on astrophysical grounds: the bounds from CAST only come into contact with our data for the largest values of N . Second, sharp edges are visible in Fig. 8: the lower edge marks the contribution to the dark matter abundance from the QCD axion itself, and therefore constitutes a lower bound on the total dark matter abundance for any given $g_{a\gamma\gamma}$. The upper edge of the distribution is set by the largest-mass axions in each theory, which in turn is dictated by the smallest volume allowed for divisors. This minimum volume was set to one in (30), but if the minimum volume were decreased by some amount, the top edge would marginally increase.

It is interesting to note the interplay of the dark matter relic densities with the θ -angle. We saw above that as the number of axions N in the effective theory increases, the stringy contribution to the $\Delta\theta$ decreases. By the same token, as N increases, the axion dark matter abundance decreases. Thus, we find a correlation between the viability of the Peccei-Quinn mechanism and the observed dark matter abundance. This hints at an anthropic solution of the strong CP problem (this possibility was also discussed in [65]).

VI. CONCLUSIONS

String theory and experiment are separated by a vast hierarchy of energy scales. The decoupling of ultraviolet physics from infrared phenomena makes science possible, but at the same time it presents an obstacle to making useful predictions from string theory. Fortunately, some mysteries of low-energy physics are sensitive to the details of the high-energy theory. The strong CP problem is a leading example: the most promising solution, the Peccei-Quinn mechanism, can be spoiled by corrections from Planck-suppressed operators. Explaining the

smallness of the neutron electric dipole moment following Peccei and Quinn thus requires understanding quantum gravity effects.

In weakly-coupled string compactifications, one can directly calculate the QCD axion potential, including the effects of the ultraviolet completion of gravity: among these the most important are D-brane instantons. We computed the leading D-brane instantons in type IIB compactifications on orientifolds of Calabi-Yau threefold hypersurfaces in toric varieties, in the geometric regime. We found that in this setting, D-brane instanton contributions to the neutron electric dipole moment fall exponentially with the fourth power of the number of axions: see (33). This realizes Polchinski's proposal [66] that the strong CP problem might admit a natural solution in topologically complex compactifications of string theory.

Assessing the QCD axion potential in these models also required understanding QCD instantons at high energies, as a question in field theory. We noted a surprising fact: although the high-energy QCD instanton contributions are small enough not to spoil the PQ mechanism in the absence of vectorlike pairs, for a low enough scale of supersymmetry breaking they make contributions to the θ -angle of a size that may be detectable with future experiments.

In the same ensemble, we carried out a preliminary analysis of the axion dark matter relic abundances, as well as effective axion-photon couplings, and compared these results to observational bounds. We found that a large fraction of models are allowed by current constraints.

We conclude that in the absence of vectorlike pairs, the strong CP problem is generically solved in this landscape of theories.

ACKNOWLEDGEMENTS

We thank Csaba Csáki, Jim Halverson, Doddy Marsh, Matt Reece, Max Ruhdorfer, and Edward Witten for discussions. We are grateful to Richard Nally for comments on a draft of this work. The research of M.D. was supported in part by the National Science Foundation under Cooperative Agreement PHY-2019786 (The NSF AI Institute for Artificial Intelligence and Fundamental Interactions). The research of M.D., N.G., and L.M. was supported in part by NSF grant PHY-1719877, that of C.L. was partially supported by DOE Grant DE-SC0013607, and that of L.M. and J.M. was supported in part by the Simons Foundation Origins of the Universe Initiative.

¹⁵ Roughly 0.1% of points in Figs. 6 and 8 have been omitted for the sake of visual clarity: in these models, the overall scale is so low that the QCD axion itself does not satisfy (35), and thus does not contribute to the dark matter density. These models have negligible dark matter abundances, and are not phenomenologically viable—in any case, they violate (7).

Appendix A: Computation of masses and decay constants

The Lagrangians considered in this paper take the form

$$\mathcal{L} \supset -\frac{1}{2} K_{AB} \partial \xi^A \partial \xi^B - \sum_{\alpha} \frac{g_A^{\alpha} \xi^A}{32\pi^2} \tilde{G}_{\alpha}^{\mu\nu} G_{\alpha\mu\nu} - V(\xi^A) \quad (\text{A1})$$

where $V(\xi^A)$ is the scalar potential induced by QCD and stringy instantons, and G denotes non-Abelian gauge field strengths (in particular, QCD). Canonically normalizing the axions ξ^A involves (1) defining a new set of axions $\phi^A = e^A_B \xi^B$ using a Cholesky decomposition $K = e^T \cdot e$ in terms of which the kinetic term is canonical, and (2) using up the residual freedom to rotate the ϕ^A with orthogonal transformations in order to diagonalize the Hessian that arises from $V(\xi^A)$. The first task is computationally straightforward, as the eigenvalues of K turn out not to be hierarchical in our models. The second step, however, is computationally expensive using brute force methods: the non-perturbative scales $\Lambda_{\vec{q}}^4$ entering $V(\xi^A)$ typically vary over many orders of magnitude, so that a direct numerical solution of the eigenvalue problem of the Hessian $\mathcal{H}_{AB} := \partial_{\xi^A} \partial_{\xi^B} V$ would require working with very high precision.

Instead, we take advantage of the fact that the non-perturbative scales are hierarchical. Consider a set of N instantons with charges \vec{q}_A , $A = 1, \dots, N$, ordered such that $S_E^{\vec{q}_A} < S_E^{\vec{q}_B}$ for all $A < B$. Without loss we assume that the kinetic matrix is the identity (in such a basis the charges \vec{q}_A take real, not integer values). The largest eigenvector of the Hessian clearly receives its dominant contribution from the instanton \vec{q}_1 with smallest action, and the largest eigenvector $\delta\phi^A$ is obtained by minimizing the distance in field space required to perturb the instanton phase $\delta\phi \cdot \vec{q}_1$ by a fixed amount, i.e. $\delta\vec{\phi} = \varphi \frac{\vec{q}_1}{\|\vec{q}_1\|}$. Along this direction in field space, the axion kinetic term and potential read

$$\mathcal{L} \supset -\frac{1}{2} (\partial\varphi)^2 - \Lambda_{\vec{q}_1}^4 \left(1 - \cos(2\pi \|\vec{q}_1\| \varphi)\right), \quad (\text{A2})$$

so one finds a decay constant¹⁶ $f_1 = (2\pi \|\vec{q}_1\|)^{-1}$, and mass $m_1^2 = \Lambda_{\vec{q}_1}^4 / f_1^2$. Below the mass scale m_1 one may integrate out the heaviest axion ϕ , and parameterize the remaining $N - 1$ axion field space as $\delta\phi^A = M^A_i \delta\hat{\phi}^i$ in terms of $N - 1$ low energy axions $\delta\hat{\phi}^i$ where the rows of M^A_i generate the kernel of the map $\vec{\phi} \mapsto \vec{q}_1 \cdot \vec{\phi}$. Again without loss one may choose M such that $M^T \cdot M =$

¹⁶ The notion of a ‘‘decay constant’’ is not well-defined for an axion potential with more than N instanton contributions, but the quantities f_i defined in this Appendix are nonetheless good proxies for the axion field range when the masses are exponentially hierarchical, as they are here. As such, we will refer to these quantities as ‘‘decay constants’’ throughout.

$\mathbb{I}_{N-1, N-1}$, and this choice is unique up to an $O(N - 1)$ transformation. The transpose M^T defines a projection $\pi : \mathbb{R}^N \rightarrow \mathbb{R}^{N-1}$, $\vec{q} \mapsto M^T \cdot \vec{q}$, and the images of the charges $\vec{q}_2, \dots, \vec{q}_N$ under π are the effective instanton charges in the low energy theory. We can repeat the same steps in the low energy theory to determine the next pair of decay constant and mass: $f_2 = (2\pi \|\pi(\vec{q}_2)\|)^{-1}$ and $m_2^2 = \Lambda_{\vec{q}_2}^4 / f_2^2$. By iterating this all the way to the smallest eigenvalue one finds that the i -th decay constant f_i is given by

$$f_i = (2\pi \|\pi_i(\vec{q}_i)\|)^{-1}, \quad (\text{A3})$$

with projections $\pi_i : \mathbb{R}^N \rightarrow \mathbb{R}^{N-i+1}$, $\vec{q} \mapsto (M_i)^T \cdot \vec{q}$ and for $i > 1$ the rows of the matrices M_i are generators of the kernel of $(\vec{q}_1, \dots, \vec{q}_{i-1})^T$, chosen such that $(M_i)^T \cdot M_i = \mathbb{I}_{N-i+1, N-i+1}$, while $M_1 := \mathbb{I}_{N, N}$. The masses are then given by $m_i^2 = \Lambda_{\vec{q}_i}^4 / f_i^2$.

The above analysis gives the correct mass spectrum in the limit $\Lambda_{\vec{q}_A} / \Lambda_{\vec{q}_{A+1}} \rightarrow \infty$. Thus, one expects $\mathcal{O}(1)$ corrections when a pair of instanton actions $S_E^{\vec{q}_A}$, $S_E^{\vec{q}_{A+1}}$ are close to each other, with $S_E^{\vec{q}_{A+1}} - S_E^{\vec{q}_A} \lesssim 1$. Indeed, in our examples this is frequently the case, but in a random sample of 100 geometries with $h^{1,1} = 10$ to 100, we found deviations in the masses to be 1% on average, with a maximum error of 80%, so we deem the above algorithm robust for our purposes.

Appendix B: Sub-leading instantons

In this Appendix, we explain why the contributions to the scalar potential studied in the main text are indeed the leading ones for the purpose of assessing the Peccei-Quinn quality problem.

a. Autochthonous divisors

The effective divisors on Calabi-Yau threefold hypersurfaces can be categorized as *inherited* or *autochthonous*. The inherited divisors, of which there are $h^{1,1} + 4$, are by definition those that descend from effective divisors of the ambient toric variety. Autochthonous divisors, on the other hand, are those that do not arise from intersections of the hypersurface with effective divisors of the ambient variety.

Enumerating the inherited divisors in a given geometry is straightforward, and in the main analysis of §IV we studied the contributions from Euclidean D3-branes wrapping inherited divisors. One might then ask whether Euclidean D3-branes wrapping autochthonous divisors could introduce a PQ quality problem in cases where Euclidean D3-branes wrapping inherited divisors do not.

Although finding all the autochthonous divisors in a general Calabi-Yau hypersurface is an open problem, a certain type of autochthonous divisor associated to factorizations of the hypersurface defining equation is easily

identified from polytope data. We studied the volumes of these so-called ‘min-face factorized’ divisors for every favorable Calabi-Yau hypersurface with $2 \leq h^{1,1} \leq 4$, evaluating the divisor volumes at the point in Kähler moduli space where all curves have volumes ≥ 1 . In this test set, we found that the smallest autochthonous divisor is on average 68 times larger than the smallest inherited divisor, and the smallest such ratio is 4.6. Min-face factorized autochthonous divisors were also studied in the compactifications of [53], where $51 \leq h^{1,1} \leq 214$, and the smallest autochthonous divisors was never less than 100 times larger than the smallest inherited divisor. Thus, barring the existence of a new class of autochthonous divisors with much smaller volumes, we expect that the PQ quality problem will be dominated by Euclidean D3-branes wrapping inherited divisors.

b. Kähler potential instantons

Thus far we have explained how we compute an upper bound on the effects of the leading *superpotential* terms. We should still consider whether PQ quality can be affected by instanton corrections to the Kähler potential that arise from Euclidean D3-branes wrapping *non-holomorphic* 4-cycles in the Calabi-Yau.

Suppose that $[\Sigma] \in H_4(X, \mathbb{Z})$ is a homology class that admits no holomorphic representative. Then a Euclidean D3-brane in the class $[\Sigma]$ will wrap the locally volume-minimizing representative Σ_{\min} of $[\Sigma]$. Because Σ_{\min} is not calibrated by the Kähler form, computing its volume is difficult. However, any class $[\Sigma]$ can be represented as a sum of classes that admit holomorphic or anti-holomorphic representatives, and so $[\Sigma]$ has a representative Σ_{\cup} that is a union of holomorphic and anti-holomorphic cycles. We call such a representative a *piecewise calibrated* representative. The piecewise-calibrated representative is generally *not* locally volume-minimizing, but one can think of the minimum-volume representative in that class as a fusing, or recombination,

of the components of the piecewise-calibrated representative. We define the recombination fraction [41]

$$\tau_{\Sigma} := \frac{\text{Vol}(\Sigma_{\cup}) - \text{Vol}(\Sigma_{\min})}{\text{Vol}(\Sigma_{\min})}. \quad (\text{B1})$$

We would like to know whether a contribution to the Kähler potential from a Euclidean D3-brane wrapping a non-holomorphic cycle Σ_{\min} in some class $[\Sigma]$ can have a larger impact on PQ quality than that of any possible superpotential term in the same compactification. This can occur only when the recombination (B1) is significant, $\tau_{\Sigma} \gtrsim 1$, as otherwise the superpotential terms corresponding to the holomorphic and antiholomorphic components of Σ_{\cup} — which were already incorporated in the analysis of §IV — will have effects that are at least comparable to the Kähler potential term.

Thus, very large recombination could spoil the PQ mechanism. However, there is no evidence for large recombination in the substantial literature on geometric measure theory: see e.g. [67]. Moreover, the investigations of [41, 42], which aimed in part to identify examples with large recombination, were able to show that (rather strong forms of) the Weak Gravity Conjecture (WGC) can imply *some* degree of recombination, but not enough to invalidate our conclusions about PQ quality. Indeed, because the axionic form of the WGC generally implies a bound $S \lesssim M_{\text{pl}}/f$, while the decay constants considered here are far below the Planck scale, it is implausible that WGC constraints alone could be stringent enough to impact PQ quality. Finally, one should recognize that if the effects of Kähler potential instantons are generally parametrically enhanced by large recombination — a conclusion that we judge to be very unlikely on the basis of the evidence given above — then this would invalidate not just our conclusions, but also much of the literature on four-dimensional theories arising in string compactifications. We leave this intriguing but implausible possibility as a question for the future.

-
- [1] NEDM collaboration, C. Abel et al., *Measurement of the permanent electric dipole moment of the neutron*, *Phys. Rev. Lett.* **124** (2020) 081803, [2001.11966].
- [2] C. A. Baker, D. D. Doyle, P. Geltenbort, K. Green, M. G. D. van der Grinten, P. G. Harris et al., *Improved experimental limit on the electric dipole moment of the neutron*, *Phys. Rev. Lett.* **97** (Sep, 2006) 131801.
- [3] J. M. Pendlebury, S. Afach, N. J. Ayres, C. A. Baker, G. Ban, G. Bison et al., *Revised experimental upper limit on the electric dipole moment of the neutron*, *Phys. Rev. D* **92** (Nov, 2015) 092003.
- [4] R. D. Peccei and H. R. Quinn, *CP Conservation in the Presence of Instantons*, *Phys. Rev. Lett.* **38** (1977) 1440–1443.
- [5] G. Grilli di Cortona, E. Hardy, J. Pardo Vega and G. Villadoro, *The QCD axion, precisely*, *JHEP* **01** (2016) 034, [1511.02867].
- [6] H. M. Georgi, L. J. Hall and M. B. Wise, *Grand Unified Models With an Automatic Peccei-Quinn Symmetry*, *Nucl. Phys. B* **192** (1981) 409–416.
- [7] R. Holman, S. D. H. Hsu, T. W. Kephart, E. W. Kolb, R. Watkins and L. M. Widrow, *Solutions to the strong CP problem in a world with gravity*, *Phys. Lett. B* **282** (1992) 132–136, [hep-ph/9203206].
- [8] S. M. Barr and D. Seckel, *Planck-scale corrections to axion models*, *Phys. Rev. D* **46** (Jul, 1992) 539–549.
- [9] M. Kamionkowski and J. March-Russell, *Planck scale physics and the Peccei-Quinn mechanism*, *Phys. Lett. B* **282** (1992) 137–141, [hep-th/9202003].
- [10] J. P. Conlon, *The QCD axion and moduli stabilisation*, *JHEP* **05** (2006) 078, [hep-th/0602233].
- [11] R. Blumenhagen, S. Moster and E. Plauschinn, *Moduli*

- Stabilisation versus Chirality for MSSM like Type IIB Orientifolds*, *JHEP* **01** (2008) 058, [0711.3389].
- [12] T. W. Grimm, M. Kerstan, E. Palti and T. Weigand, *On Fluxed Instantons and Moduli Stabilisation in IIB Orientifolds and F-theory*, *Phys. Rev. D* **84** (2011) 066001, [1105.3193].
- [13] I. Broeckel, M. Cicoli, A. Maharana, K. Singh and K. Sinha, *Moduli stabilisation and the statistics of axion physics in the landscape*, *JHEP* **08** (2021) 059, [2105.02889].
- [14] M. Bies, M. Cvetič, R. Donagi, M. Liu and M. Ong, *Root Bundles and Towards Exact Matter Spectra of F-theory MSSMs*, 2102.10115.
- [15] J. Preskill, M. B. Wise and F. Wilczek, *Cosmology of the Invisible Axion*, *Phys. Lett. B* **120** (1983) 127–132.
- [16] L. F. Abbott and P. Sikivie, *A Cosmological Bound on the Invisible Axion*, *Phys. Lett. B* **120** (1983) 133–136.
- [17] M. Dine and W. Fischler, *The Not So Harmless Axion*, *Phys. Lett. B* **120** (1983) 137–141.
- [18] M. Dine and N. Seiberg, *String Theory and the Strong CP Problem*, *Nucl. Phys. B* **273** (1986) 109–124.
- [19] N. Arkani-Hamed, S. Dimopoulos and S. Kachru, *Predictive landscapes and new physics at a TeV*, hep-th/0501082.
- [20] G. Dvali, *Black Holes and Large N Species Solution to the Hierarchy Problem*, *Fortsch. Phys.* **58** (2010) 528–536, [0706.2050].
- [21] G. Dvali and M. Redi, *Black Hole Bound on the Number of Species and Quantum Gravity at LHC*, *Phys. Rev. D* **77** (2008) 045027, [0710.4344].
- [22] M. Demirtas, C. Long, L. McAllister and M. Stillman, *The Kreuzer-Skarke Axiverse*, *JHEP* **04** (2020) 138, [1808.01282].
- [23] C. Vafa and E. Witten, *Restrictions on Symmetry Breaking in Vector-Like Gauge Theories*, *Nucl. Phys. B* **234** (1984) 173–188.
- [24] C. Vafa and E. Witten, *Parity Conservation in QCD*, *Phys. Rev. Lett.* **53** (1984) 535.
- [25] H. Georgi and L. Randall, *Flavor Conserving CP Violation in Invisible Axion Models*, *Nucl. Phys. B* **276** (1986) 241–252.
- [26] B. Holdom, *Strong QCD at High-energies and a Heavy Axion*, *Phys. Lett. B* **154** (1985) 316.
- [27] J. M. Flynn and L. Randall, *A Computation of the Small Instanton Contribution to the Axion Potential*, *Nucl. Phys. B* **293** (1987) 731–739.
- [28] K. Choi and H. D. Kim, *Small instanton contribution to the axion potential in supersymmetric models*, *Phys. Rev. D* **59** (1999) 072001, [hep-ph/9809286].
- [29] C. Csáki, M. Ruhdorfer and Y. Shirman, *UV Sensitivity of the Axion Mass from Instantons in Partially Broken Gauge Groups*, *JHEP* **04** (2020) 031, [1912.02197].
- [30] T. Gherghetta, V. V. Khoze, A. Pomarol and Y. Shirman, *The Axion Mass from 5D Small Instantons*, *JHEP* **03** (2020) 063, [2001.05610].
- [31] R. Kitano and W. Yin, *Strong CP problem and axion dark matter with small instantons*, *JHEP* **07** (2021) 078, [2103.08598].
- [32] A. Hook, *TASI Lectures on the Strong CP Problem and Axions*, *PoS TASI2018* (2019) 004, [1812.02669].
- [33] W. Buchmuller and D. Wyler, *Effective Lagrangian Analysis of New Interactions and Flavor Conservation*, *Nucl. Phys. B* **268** (1986) 621–653.
- [34] A. Canepa, *Searches for Supersymmetry at the Large Hadron Collider*, *Rev. Phys.* **4** (2019) 100033.
- [35] M. Cvetič, J. Halverson and P. Langacker, *Implications of String Constraints for Exotic Matter and Z' s Beyond the Standard Model*, *JHEP* **11** (2011) 058, [1108.5187].
- [36] J. Halverson, *Anomaly Nucleation Constrains SU(2) Gauge Theories*, *Phys. Rev. Lett.* **111** (2013) 261601, [1310.1091].
- [37] P. Svrček and E. Witten, *Axions In String Theory*, *JHEP* **06** (2006) 051, [hep-th/0605206].
- [38] M. Cicoli, M. Goodsell and A. Ringwald, *The type IIB string axiverse and its low-energy phenomenology*, *JHEP* **10** (2012) 146, [1206.0819].
- [39] T. C. Bachlechner, K. Eckerle, O. Janssen and M. Kleban, *The Axidental Universe*, 1902.05952.
- [40] M. Cicoli, V. Guidetti, N. Righi and A. Westphal, *Fuzzy Dark Matter Candidates from String Theory*, 2110.02964.
- [41] M. Demirtas, C. Long, L. McAllister and M. Stillman, *Minimal Surfaces and Weak Gravity*, *JHEP* **03** (2020) 021, [1906.08262].
- [42] C. Long, A. Sheshmani, C. Vafa and S.-T. Yau, *Non-Holomorphic Cycles and Non-BPS Black Branes*, 2104.06420.
- [43] M. Kreuzer and H. Skarke, *Complete classification of reflexive polyhedra in four-dimensions*, *Adv. Theor. Math. Phys.* **4** (2002) 1209–1230, [hep-th/0002240].
- [44] V. V. Batyrev, *Dual Polyhedra and Mirror Symmetry for Calabi-Yau Hypersurfaces in Toric Varieties*, *J. Alg. Geom.* (1996) 493–535, [alg-geom/9310003].
- [45] M. Demirtas, L. McAllister and A. Rios-Tascon, *Bounding the Kreuzer-Skarke Landscape*, 2008.01730.
- [46] M. Demirtas, L. McAllister and A. Rios-Tascon, *CYTools: A Software Package for Analyzing Calabi-Yau Manifolds*, to appear.
- [47] J. Halverson, C. Long, B. Nelson and G. Salinas, *Towards string theory expectations for photon couplings to axionlike particles*, *Phys. Rev. D* **100** (2019) 106010, [1909.05257].
- [48] V. M. Mehta, M. Demirtas, C. Long, D. J. E. Marsh, L. McAllister and M. J. Stott, *Superradiance Exclusions in the Landscape of Type IIB String Theory*, 2011.08693.
- [49] V. M. Mehta, M. Demirtas, C. Long, D. J. E. Marsh, L. McAllister and M. J. Stott, *Superradiance in string theory*, *JCAP* **07** (2021) 033, [2103.06812].
- [50] G. Veneziano and S. Yankielowicz, *An Effective Lagrangian for the Pure N=1 Supersymmetric Yang-Mills Theory*, *Phys. Lett. B* **113** (1982) 231.
- [51] I. Affleck, M. Dine and N. Seiberg, *Dynamical Supersymmetry Breaking in Supersymmetric QCD*, *Nucl. Phys. B* **241** (1984) 493–534.
- [52] F. Denef, M. R. Douglas, B. Florea, A. Grassi and S. Kachru, *Fixing all moduli in a simple f-theory compactification*, *Adv. Theor. Math. Phys.* **9** (2005) 861–929, [hep-th/0503124].
- [53] M. Demirtas, M. Kim, L. McAllister, J. Moritz and A. Rios-Tascon, *Small Cosmological Constants in String Theory*, 2107.09064.
- [54] D. J. E. Marsh, *Axion Cosmology*, *Phys. Rept.* **643** (2016) 1–79, [1510.07633].
- [55] D. Cyncynates, T. Giurgica-Tiron, O. Simon and J. O. Thompson, *Friendship in the Axiverse: Late-time direct and astrophysical signatures of early-time nonlinear axion dynamics*, 2109.09755.

- [56] F. Chadha-Day, *Axion-like particle oscillations*, *JCAP* **01** (2022) 013, [2107.12813].
- [57] D. J. E. Marsh and P. G. Ferreira, *Ultra-Light Scalar Fields and the Growth of Structure in the Universe*, *Phys. Rev. D* **82** (2010) 103528, [1009.3501].
- [58] CAST collaboration, V. Anastassopoulos et al., *New CAST Limit on the Axion-Photon Interaction*, *Nature Phys.* **13** (2017) 584–590, [1705.02290].
- [59] D. Wouters and P. Brun, *Constraints on Axion-like Particles from X-Ray Observations of the Hydra Galaxy Cluster*, *Astrophys. J.* **772** (2013) 44, [1304.0989].
- [60] M. Berg, J. P. Conlon, F. Day, N. Jennings, S. Krippendorff, A. J. Powell et al., *Constraints on Axion-Like Particles from X-ray Observations of NGC1275*, *Astrophys. J.* **847** (2017) 101, [1605.01043].
- [61] M. C. D. Marsh, H. R. Russell, A. C. Fabian, B. P. McNamara, P. Nulsen and C. S. Reynolds, *A New Bound on Axion-Like Particles*, *JCAP* **12** (2017) 036, [1703.07354].
- [62] J. P. Conlon, F. Day, N. Jennings, S. Krippendorff and M. Rummel, *Constraints on Axion-Like Particles from Non-Observation of Spectral Modulations for X-ray Point Sources*, *JCAP* **07** (2017) 005, [1704.05256].
- [63] L. Chen and J. P. Conlon, *Constraints on massive axion-like particles from X-ray observations of NGC 1275*, *Mon. Not. Roy. Astron. Soc.* **479** (2018) 2243–2248, [1712.08313].
- [64] C. S. Reynolds, M. C. D. Marsh, H. R. Russell, A. C. Fabian, R. Smith, F. Tombesi et al., *Astrophysical limits on very light axion-like particles from Chandra grating spectroscopy of NGC 1275*, 1907.05475.
- [65] M. Dine, L. Stephenson Haskins, L. Ubaldi and D. Xu, *Some Remarks on Anthropic Approaches to the Strong CP Problem*, *JHEP* **05** (2018) 171, [1801.03466].
- [66] J. Polchinski, *The Cosmological Constant and the String Landscape*, in *23rd Solvay Conference in Physics: The Quantum Structure of Space and Time*, pp. 216–236, 3, 2006, hep-th/0603249.
- [67] M. Micalef and J. Wolfson, *Area minimizers in a $K3$ surface and holomorphicity*, *GAGA Geometric And Functional Analysis* **16** (01, 2006) 437–452.