

High Quality QCD Axion and the LHC

Anson Hook,[†] Soubhik Kumar,^{Ⓜ,*} Zhen Liu,^{Ⓜ,‡} and Raman Sundrum[§]
*Maryland Center for Fundamental Physics, Department of Physics, University of Maryland,
 College Park, Maryland 20742, USA*

 (Received 23 December 2019; accepted 29 April 2020; published 1 June 2020)

The QCD axion provides an elegant solution to the strong CP problem. While the minimal realization is vulnerable to the so-called “axion quality problem,” we will consider a more robust realization in the presence of a mirror sector related to the standard model by a (softly broken) \mathbb{Z}_2 symmetry. We point out that the resulting “heavy” axion, while satisfying all theoretical and observational constraints, has a large and uncharted parameter space, which allows it to be probed at the LHC as a long-lived particle (LLP). The small defining axionic coupling to gluons results in a challenging hadronic decay signal which we argue can be distinguished against the background in such a long-lived regime, and yet, the same coupling allows for sufficient production at the hadron colliders thanks to the large gluon-parton luminosity. Our study opens up a new window towards accelerator observable axions and, more generally, singly produced LLPs.

DOI: 10.1103/PhysRevLett.124.221801

Introduction.—The strong CP problem is the puzzle of why the strong interactions are CP symmetric, even though the standard model (SM) as a whole is not. Technically, the question centers on the vanishingly small value of the one CP -violating coupling of QCD, θ . An elegant solution is provided by elevating θ to a fully dynamical pseudoscalar “axion” field. If the axion gets its potential entirely through QCD effects, then remarkably its ground state automatically corresponds to $\theta = 0$ [1–4].

Despite this bottom-up simplicity, the QCD axion mechanism has a top-down flaw: the axion quality problem [5–8]. This arises because there can be other UV contributions to the axion potential that can push the minimum away from $\theta = 0$. We consider a mirror sector in the UV related to the SM by a \mathbb{Z}_2 symmetry, generalizing on Refs. [9–13], coupled to the same axion, such that its contribution to the axion potential is much larger than QCD’s but aligned with it in having its minimum at $\theta = 0$. This results in an axion mechanism which is vastly more robust against other uncorrelated UV effects. See Refs. [14–17] (also recent work Refs. [18–20]) for special UV structures that mitigate the quality problem, and Refs. [21–28] for alternate mechanisms for having a heavy QCD axion.

Our general solution to the quality problem places us in a very different and interesting region in axion mass-coupling parameter space (for reviews of the standard regimes, see Refs. [29,30]), in which the axion can be

probed as a quantum particle at cutting-edge collider experiments. While there are a number of axion(-like) particle collider studies [31–65], we will show that this high quality axion search presents an experimentally novel but challenging target for singly produced long-lived particles (LLPs) at the LHC main detectors.

Axion mechanism and the quality problem.—The QCD axion field $a(x)$ is coupled to QCD by promoting $\bar{\theta} \rightarrow \bar{\theta} + a(x)/f_a$:

$$\mathcal{L} \supset \frac{\alpha_3}{8\pi} \left(\bar{\theta} + \frac{a(x)}{f_a} \right) G_{\mu\nu}^a \tilde{G}^{a,\mu\nu}. \quad (1)$$

In the absence of the axion, $\bar{\theta}$ represents the CP -odd gauge invariant QCD coupling, constrained by bounds on the neutron electric dipole moment (EDM) to be $|\bar{\theta}| < 10^{-10}$ [66]. In the above, $\alpha_3 = g_s^2/4\pi$ is given in terms of the QCD gauge coupling g_s ; \tilde{G} denotes the dual of the gluon field strength, $\tilde{G}^{a,\mu\nu} = \frac{1}{2} \epsilon^{\mu\nu\rho\sigma} G_{\rho\sigma}^a$ with $\epsilon^{0123} = 1$; and f_a denotes the axion “decay constant.” The parameter $\bar{\theta}$ is the *effective* θ parameter obtained after diagonalizing the Yukawa matrices via chiral rotations, and is given by $\bar{\theta} \equiv \theta + \arg \det(Y_u Y_d)$, where $Y_{u(d)}$ is the up (down) type Yukawa matrix with complex entries and θ is a *bare* Lagrangian parameter.

The nonperturbative QCD axion potential resulting from Eq. (1) can be calculated using chiral perturbation theory [3,67],

$$\mathcal{V} \approx -m_\pi^2 f_\pi^2 \sqrt{1 - \frac{4m_u m_d}{(m_u + m_d)^2} \sin^2 \left(\frac{a(x)}{2f_a} + \frac{\bar{\theta}}{2} \right)}, \quad (2)$$

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International 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.

where m_π and f_π are, respectively, the mass and the decay constant of the pion, $m_{u(d)}$ is the up (down) quark mass. Given the potential in Eq. (2) (and its refinements), the axion acquires a vacuum expectation value (VEV), $\langle a \rangle = -f_a \bar{\theta}$. Plugging this into Eq. (1) we see at low energies that the CP violation in QCD is eliminated—solving the strong CP problem.

Clearly, if there is any other contribution to the axion potential from beyond QCD, the resulting axion ground state need no longer screen the QCD CP violation. One can realize the axion as a Nambu-Goldstone boson of a $U(1)$ Peccei-Quinn (PQ) symmetry, $e^{ia/f_a} \rightarrow e^{i\phi} e^{ia/f_a}$, which would forbid any axion potential. But such a symmetry cannot be exact because it is broken by QCD (chiral anomaly) effects, and further in the far UV quantum gravity is expected to explicitly break all global symmetries [68,69]. At best, PQ must be an accidental symmetry of the leading couplings. For example, the UV violation of PQ symmetry may take the form of a higher-dimensional composite operator \mathcal{O} with scaling dimension $\Delta > 4$ and PQ charge q ,

$$\mathcal{L} \supset \frac{\mathcal{O}}{M_p^{\Delta-4}} \underset{\text{low energies}}{\sim} \frac{f_a^\Delta}{M_p^{\Delta-4}} e^{iqa/f_a}, \quad (3)$$

where $M_p \sim 10^{19}$ GeV is the Planck scale. Naively, one would expect that Planck-suppressed PQ violation would have negligible effects on the axion mechanism. However, given the experimental or astrophysical constraint $f_a > 10^9$ GeV [30] and the very delicate QCD potential of Eq. (2), the axion mechanism is spoiled (for $q \sim 1$) out to very high scaling dimension unless $\Delta \geq 9$. This extreme fragility of the axion mechanism is the so-called “quality problem” [5–8].

Model construction.—While it is possible, but demanding, that the UV structure does strongly suppress PQ violation at this level, in this Letter, we study a different approach: we strengthen the IR axion mechanism itself. We will consider a \mathbb{Z}_2 symmetry that exchanges the fields of the SM and matching fields of a mirror sector. This replication includes the SM gauge structure, so that the entire mirror sector carries no SM charges and vice versa. All the marginal (dimensionless) couplings of the two sectors are identical, including $\bar{\theta}$. However \mathbb{Z}_2 is softly broken in the one relevant operator by having two distinct tachyonic Higgs mass terms, $\mu^2 H^\dagger H$ and $\mu'^2 H'^\dagger H'$. We will consider $M_p^2 \gg \mu'^2 \gg \mu^2$. One can view $\mu'^2 - \mu^2$ as arising from the VEV of a \mathbb{Z}_2 -odd scalar coupled to the Higgs fields in the far UV. It is then plausible that the \mathbb{Z}_2 breaking in the marginal couplings is as small as $\sim \mu'^2/M_p^2$.

We now introduce a \mathbb{Z}_2 -invariant QCD axion coupling,

$$\frac{\alpha_3}{8\pi} \left(\frac{a}{f_a} + \bar{\theta} \right) (G\tilde{G} + G'\tilde{G}'). \quad (4)$$

The equality of the strong couplings indicated above is true in the far UV, by \mathbb{Z}_2 , but can run differently below $\sim \mu'$. For $\mu' > 100$ TeV, $\Lambda_{\text{QCD}} \ll \Lambda'_{\text{QCD}} < m_{q'}$ for all mirror quarks q' . We can estimate the strong coupling scale Λ'_{QCD} given Λ_{QCD} using the 1-loop renormalization group (RG). The differential running at 1 loop depends on the quark and mirror quark masses in terms of μ and μ' but is insensitive to new model-dependent thresholds involving colored degrees of freedom needed to UV complete the nonrenormalizable axionic couplings, which do not get mass through electro-weak symmetry breaking.

In the above regime the nonperturbative QCD' (pure glue) contribution to the axion potential near its minimum is given by lattice computation [70], and continuum ($\overline{\text{MS}}$) matching [71]:

$$\mathcal{V}' = 0.3 (\text{few} \times \Lambda'_{\text{QCD}})^4 \left(\frac{a}{f_a} + \bar{\theta} \right)^2 + \mathcal{O} \left[\left(\frac{a}{f_a} + \bar{\theta} \right)^4 \right], \quad (5)$$

where “few” $\sim \alpha_3^{-0.4}$ is in detail the model-dependent conversion factor between the 1 loop and 2 loop estimates of the strong coupling scales.

Considering Eq. (4), this immediately shows that the single QCD axion a solves the strong CP problems of both the sectors at the same time by having the VEV $\langle a \rangle = -f_a \bar{\theta}$. Although the two values of $\bar{\theta}$ are identical for the two sectors in the UV, the breaking of the \mathbb{Z}_2 symmetry can make the two $\bar{\theta}$'s different below μ' . However, RG running of $\bar{\theta}$ occurs at seven loops, and contributions from threshold corrections arise at four loops [72]. Thus, both of these effects, arising from renormalizable operators in SM and SM', are too small to be significant, even given the tight neutron EDM constraints.

However, higher dimensional operators, suppressed by the Planck scale, can make the two angles different. For example, the interactions

$$\frac{\alpha_3}{8\pi} \left(\frac{H^\dagger H}{M_p^2} G\tilde{G} + \frac{H'^\dagger H'}{M_p^2} G'\tilde{G}' \right) \quad (6)$$

can give $\bar{\theta} \neq \bar{\theta}'$ upon the breaking of the \mathbb{Z}_2 symmetry. We need the splitting $|\bar{\theta} - \bar{\theta}'| < 10^{-10}$ in order to maintain the axion mechanism to within current neutron EDM constraint. This implies $\mu' < 10^{14}$ GeV, and thus there is a maximal amount by which a can be made heavier. For further theoretical considerations see the Supplemental Material [73], which contains Refs. [74–76].

From Eq. (5), we see that the resulting axion mass is much larger than in the SM alone (for a given f_a) so that it can be heavier than Λ_{QCD} . This significantly weakens the existing experimental constraints and allows stronger couplings, $1/f_a$, to the SM. The raising of the axion mass and lowering of f_a clearly reduces the severity of the quality problem. This opens up a strongly motivated and

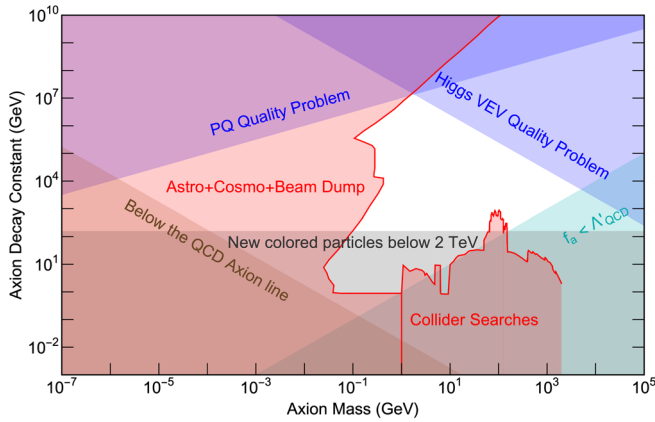


FIG. 1. The preferred model parameter regions for our high quality axion model. We require that no new colored particles exist with a mass below 2 TeV; higher dimensional operators involving the axion or the Higgs do not reintroduce the strong CP problem, as well as several astrophysical, cosmological, and collider constraints.

new experimentally testable regime for the QCD axion, which we identify now.

Constraints on parameter space.—In Fig. 1 we show the preferred parameter space for our model. We begin with the quality problem. We will choose as a benchmark a composite axion model for which PQ symmetry holds at the renormalizable level but can be violated at $\Delta = 6$. Given Eq. (3), this reintroduces the strong CP problem in the region labeled “PQ Quality Problem” in Fig. 1. We cannot populate the area labeled “Below the QCD Axion line” as our mechanism can only make the axion heavier, and not lighter. In the area labeled “Higgs VEV Quality Problem,” $\langle H' \rangle \sim \mu' > 10^{14}$ GeV and Planck suppressed operators spoil the axion mechanism, as explained around Eq. (6). Our effective field theory is only valid if $f_a > \Lambda_{\text{QCD}}$ which excludes the region shaded in cyan in Fig. 1. We are being agnostic about the origin of the axion coupling to QCD. Typically, the coupling is generated by integrating out colored fundamental fermions who get a mass $y f_a e^{ia/f_a}$ [77]. Requiring that the Yukawa coupling y of these fermions is smaller than 4π , we have new colored particles below $4\pi f_a$. Requiring that these colored fermions satisfy LHC constraints [78–81] and are heavier than ~ 2 TeV disfavors the region shown in Fig. 1. The regions labeled “Astro + Cosmo + Beam Dump” and “Collider Searches” are ruled out due to a variety of supernova, stellar cooling, beam dump, collider, and cosmology constraints. All these constraints, along with original references, can be found in the Supplemental Material [73] and references therein [36,82–94]. We see then that the most favored region is given by $m_a \sim \text{GeV-TeV}$ and $f_a \sim \text{TeV-PeV}$, ripe for collider exploration.

Phenomenology.—We present a search strategy and discuss its feasibility for massive axions in the GeV to tens of GeV range, with decay constants ranging between

100 TeV and PeV, thereby covering a sizable portion of the open regime seen in Fig. 1. We note here that the axion predominantly decays into gluons which give rise to hadronic final states. Hence, with the production mode we consider below, the signal will be a displaced jet recoiling against prompt jet(s), where the displaced jet should contain at least three tracks from axion decay with the properties described below. The presence of a displaced vertex gives us a powerful discriminator against hadronic backgrounds while still having a sizable production rate. The existence of such a parameter regime, where the signal is produced with both a significant rate and with significant displacement, is nontrivial for the GeV-scale axion, or, in general, for any singly produced long-lived particles, given that the same small coupling controls the production rate and upper limit of the proper lifetime. The lever arm that offsets the small production coupling is provided by the immense gluon-parton luminosity at the LHC and other proton accelerator experiments.

The production rate of the axion at the 13 TeV LHC is

$$\sigma(pp \rightarrow a + X; H_T > 100 \text{ GeV}) \simeq \left(\frac{100 \text{ TeV}}{f_a} \right)^2 \times 0.7 \text{ fb}.$$

Here we have imposed a cut on the sum of transverse energy of hadronic activity, $H_T > 100$ GeV and hence, the axion mass ($\lesssim 20$ GeV) does not significantly affect the cross section. If one instead requires a leading (prompt) jet minimal p_T cut of 30 GeV, the cross section increases by around a factor of 3. For further details of the cross section calculation, done using MadGraph5@NLO [95], PYTHIA8 [96,97], and DELPHES3 [98], see the Supplemental Material [73]. The lifetime of the axion can be approximated by

$$c\tau \simeq 0.8 \left(\frac{f_a}{100 \text{ TeV}} \right)^2 \left(\frac{2 \text{ GeV}}{m_a} \right)^3 \text{ mm}. \quad (7)$$

A good signal trigger efficiency is critical given the above cross section, which can be achieved using the displaced track trigger discussed in Refs. [99,100]. We follow this construction with conservative modifications to accommodate our signal. In detail, we require (i) At least three $p_T > 2$ GeV tracks within an Level-1 (L1) trigger jet, (ii) Among the above tracks, at least three of them have the transverse impact parameter $d_0 > 1$ mm, (iii) The pseudorapidity of the tracks to be $|\eta| < 2.4$, (iv) The event has $H_T > 100$ GeV. Additionally, in our analysis, we require decay location of axion candidates in the transverse plane with displacement $d_T < 35$ cm to have more than five hits in the trackers. This requirement of the number of hits for a displaced track has already been folded into to the L1 trigger rates dominated by the “fake-track” backgrounds, coming from misconnections of the tracker hits or instrumental noise. These requirements are sufficient to pass the

L1 trigger with affordable rates below $\lesssim 10$ kHz. Such an L1 rate is dominated by backgrounds from fake tracks as per the displaced track trigger in Refs. [99,100]. The requirement of three or more tracks can be passed quite easily for axion masses \gtrsim few GeV. Other SM background from metastable SM hadron decays can be effectively vetoed with the number of displaced tracks and vertexing requirements; for more detailed considerations, see the Supplemental Material [73] including Refs. [101–104].

The fake-track background is one of the most crucial for displaced vertex searches at the LHC at every stage of experimental analysis and is very large. We will use conservative estimates to suggest that this background can be cut to a small enough size to reveal the signal. Our estimates are not, however, a substitute for a full experimental analysis, which we hope to inspire. The key feature of these fake tracks is that they allow for a much larger reconstructed impact parameter than the SM background. Empirically, one can model them as tracks with flat distributions in the finite range in the following dimensions: track impact parameter ($|d_0| < 15$ cm), track curvature [$1/R \propto q/p_T < 1/(1.8$ m)], track eta $|\eta| < 2.4$, track time t_0 ($|t_0| < 6$ ns), and track z coordinate ($|z_0| < 15$ cm). These definitions of L1 fake-track parameters are mainly based upon Ref. [105]. We note that these assumptions are rough and can vary based upon detector performance and collider environments, however, they can serve as a good benchmark to start the discussion of such LLP search strategies.

Nevertheless, the fake tracks that pass the above L1 trigger constitute huge background, amounting to about $10 \text{ kHz} \times 10^8 \text{ sec} = 10^{12}$ triples of fake tracks (in jets) at the HL-LHC. Combinations of SM displaced tracks and fake tracks are expected to give a subdominant background; for more discussions, see the Supplemental Material [73]. At higher level triggers and in the analysis, one needs to suppress the background much further, to below a Hz while maintaining a high signal efficiency. Beyond all existing displaced trigger studies, we further demonstrate that it is possible to suppress these backgrounds using a 2D-4D displaced vertexing selection at a high-level trigger. We first solve for a 2D vertex by finding the best-fit point that minimizes the distances between the vertex and the tracks. Then we construct the 4D vertex of the system by extrapolating the 2D point in the transverse plane to the z direction and time direction by propagating the tracks. The time direction is defined as the timing information of each track with respect to the reference time, e.g., the beam-crossing time [106].

Our 2D-4D displaced vertexing selection is defined as follows: (1) The 2D tracks fit a common vertex with standard deviation $\Delta d_T < 1$ cm; (2) The 2D common vertex has a minimal distance to the interaction point of 0.5 cm and maximal distance of 35 cm, $0.5 \text{ cm} < d_T < 35$ cm; (3) The 2D common vertex is significantly displaced away from the

interaction point, $d_T/\Delta d_T > 5$; (4) The corresponding 4D vertex has a standard deviation in the z direction $\Delta d_z < 5$ cm; (5) The corresponding 4D vertex has a z direction location $d_z < 20$ cm; (6) The corresponding 4D vertex has a standard deviation in time $\Delta d_t < 500$ ps; (7) The corresponding 4D vertex has a time $d_t < 1000$ ps; (8) The tracks are within 0.4 in pseudorapidity of the reconstructed displaced jet direction $|\eta_i - \eta_V| < 0.4$ for all the three tracks; (9) The tracks are within 0.4 in azimuthal angle of the reconstructed displaced jet direction $|\phi_i - \phi_V| < 0.4$ for all the three tracks, where η_i (ϕ_i) and η_V (ϕ_V) are the pseudorapidity (azimuthal angle) of the displaced tracks and primary vertex, respectively.

Following the empirical model of fake-track distribution discussed above, we find that the combination of the transverse plane vertex fitting (Cut1 + Cut2 + Cut3) provides a suppression factor of 8.2×10^{-2} . Note, cuts 2 and 3 remove prompt SM backgrounds and processes involving shorter decay lengths, while Cut1 removes fake tracks. The combination of z -direction consistency (Cut4 + Cut5) and t -direction consistency (Cut6 + Cut7) provides 4.9×10^{-2} and 3.0×10^{-3} background suppression, respectively. Furthermore, the requirement for the displaced tracks pointing back to the primary vertex (Cut8 + Cut9) provides 4.9×10^{-4} suppression. After taking into account the correlations between the selection cuts (see the Supplemental Material [73]), the resulting overall suppression factor of the fake-track background from this 2D-4D vertex fitting procedure is 2.9×10^{-9} . This means the background is reduced to $10^{12} \times 2.9 \times 10^{-9} = 2900$ events. A crucial consideration on top of the above background estimation is that it is using outer layers of the tracking information only. For the signal, there would be consistent energy deposition in the electromagnetic calorimeter and hadronic calorimeter, as well as inner tracker information, which will improve the spatial resolution of the displaced tracks and constitute a powerful consistency check and discrimination of signal from fake-track background. We further assume, but leave for the experiments to demonstrate, that the matching of the information between different subdetectors for all the tracks and as well the neutral hadrons, should be able to at least provide 1 order of magnitude fake-track background suppression per track, reducing our background estimate to $2900 \times (10^{-1})^3 \simeq 3$ events. Hence, it is plausible that the fake-track background can be suppressed to negligible levels.

Depending on the detector performance at the HL-LHC, one can have 10–30 fake tracks per collision and the ranges in the fake-track distribution model may vary. The study in Ref. [100] showed that by requiring two fake tracks, one can have $\sim 10^{-1}$ background rate suppression and the H_T cut provides $\sim 10^{-2}$ suppression, so the overall rate at L1 is around 10 kHz. Here the fake tracks are those which pass the L1 track fitting procedure as described in Ref. [100], which requires the consistency between the track hits in the

p_T modules for the CMS detector. Our evaluation here conservatively assumes that the same level of suppression can be achieved by requiring three fake tracks. As shown in the discussion of the cross section in the Supplemental Material [73], such an H_T cut reduces the inclusive cross section by more than 2 orders of magnitude for the axion mass regions of interest. To show what one can achieve with a slightly less conservative trigger consideration, in the next section, we also consider a trigger with a leading-jet minimal p_T of 30 GeV, instead of the minimal H_T of 100 GeV, plus three displaced tracks. We assume the same level of the L1 rate and background can be maintained using advanced trigger developments such as matching information between different subdetectors.

Results.—After these considerations where we show that it is at least plausible to suppress the background to $\mathcal{O}(10)$ level using vertex fitting and information matching between detectors, we give the estimated sensitivities for our signals at the HL-LHC with 3000 fb^{-1} integrated luminosity. We show in Fig. 2 the reach of the proposed search with 3 and 10 signal events. These correspond to the required number of signal events for 95% C.L. exclusions for the two assumptions of 0 and 25 background events, respectively. As can be seen from Fig. 2 the importance of experiments in actually reducing the background to $\mathcal{O}(10)$ level is crucial to conduct such a high-quality axion search.

There are several unique features of the coverage plot in Fig. 2, which are different from the analogous results involving more common LLP searches where LLPs are pair produced. Understanding this will help in designing and optimizing future searches for GeV scale axions and, more generally, singly produced LLPs. First, the coverage has a strong dependence on the number of signal events needed, which is clearly shown by comparing the coverage between the shaded regions with different colors. The reason behind

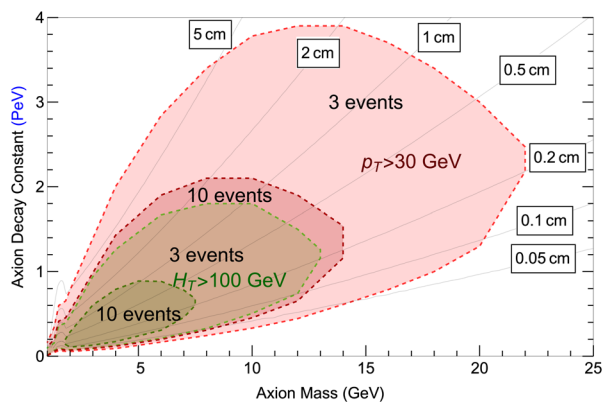


FIG. 2. The signal coverage in the m_a and f_a plane where the red and green contours correspond to requiring a leading jet with $p_T > 30$ and $H_T > 100$ GeV, respectively. For both the red and green contours, we show 10 and 3 signal event regions that can be covered following the present work. The reference lines for various ALP proper lifetimes are shown in solid grey.

the strong dependence comes from our trigger requirement $H_T > 100$ GeV or leading-jet $p_T > 30$ GeV. Given that the axion mass we are probing is much smaller than these scales, the production cross section remains approximately constant, as discussed around Eq. (7) and shown clearly in the Supplemental Material [73]. Second, the lower edge (low mass and low decay constant) of the coverage is determined by the minimal displacement requirement, below which the probability becomes too low for a GeV axion to be sufficiently displaced. Third, unlike most LLPs that are produced by stronger interactions than those involved in their decays, the upper edge (high mass and high decay constants) for our search is limited by the production rate of the axion.

Discussion and outlook.—The quest for axions is pressing. We have put forward a general theoretical structure involving a mirror sector in which a true QCD axion solves the strong CP problem while being very robust against the axion quality problem. We find that $\sim\text{GeV}$ axions with $\sim\text{PeV}$ decay constants lie at the heart of the motivated parameter space (see Fig. 1).

For $\sim\text{GeV}$ axions with $\sim\text{PeV}$ decay constants, the axion can be produced and detected at the LHC as a long-lived particle. This is a challenging long-lived particle search at the LHC. The production rate is highly suppressed by the same small coupling that leads to the displaced decay, implying that for reasonable rates, most decays will take place inside the LHC main detectors. Furthermore, there is only one low-mass displaced vertex in the event, while most other searches are for pair production of massive long-lived particles. We explored the dominant background of fake tracks and proposed a three-displaced-track strategy with 2D-4D displaced vertex reconstruction, demonstrating that the background can be feasibly suppressed. We believe that this makes the case for experimental exploration. An LLP search along the lines described in this Letter will require creative designs of the triggers and analysis at the LHC main detectors.

The codes for this Letter are available at Axion@LHC [107].

We would like to thank Matthew Daniel Citron, Jared Evans, Yuri Gershtein, Simon Knapen, and Diego Redigolo for very useful comments on the draft. We would also like to thank Prateek Agrawal, Evan Berkowitz, Zohreh Davoudi, and Simone Pagan Griso for helpful discussions. This research was supported in part by the NSF Grants No. PHY-1620074, No. PHY-1914480, and No. PHY-1914731, and by the Maryland Center for Fundamental Physics (MCFP). A. H., Z. L., and R. S. acknowledge the hospitality of the Kavli Institute for Theoretical Physics, UC Santa Barbara, during the “Origin of the Vacuum Energy and Electroweak Scales” workshop, and the support by the NSF Grant No. PHY-174958. A. H. and Z. L. would also like to thank Pittsburgh Particle Physics Astrophysics

and Cosmology Center (PITT-PACC), The Munich Institute for Astro- and Particle Physics (MIAPP), and Aspen Center for Physics (supported by NSF Grant No. PHY-1607611) for support from their programs and providing the environment for collaboration during various stages of this work.

*Corresponding author.

soubhik@terpmail.umd.edu

[†]hook@umd.edu

[‡]zliuphys@umd.edu

[§]raman@umd.edu

- [1] R. D. Peccei and H. R. Quinn, *Phys. Rev. Lett.* **38**, 1440 (1977).
- [2] R. D. Peccei and H. R. Quinn, *Phys. Rev. D* **16**, 1791 (1977).
- [3] S. Weinberg, *Phys. Rev. Lett.* **40**, 223 (1978).
- [4] F. Wilczek, *Phys. Rev. Lett.* **40**, 279 (1978).
- [5] M. Kamionkowski and J. March-Russell, *Phys. Lett. B* **282**, 137 (1992).
- [6] S. M. Barr and D. Seckel, *Phys. Rev. D* **46**, 539 (1992).
- [7] S. Ghigna, M. Lusignoli, and M. Roncadelli, *Phys. Lett. B* **283**, 278 (1992).
- [8] R. Holman, S. D. H. Hsu, T. W. Kephart, E. W. Kolb, R. Watkins, and L. M. Widrow, *Phys. Lett. B* **282**, 132 (1992).
- [9] V. A. Rubakov, *JETP Lett.* **65**, 621 (1997).
- [10] Z. Berezhiani, L. Gianfagna, and M. Giannotti, *Phys. Lett. B* **500**, 286 (2001).
- [11] A. Hook, *Phys. Rev. Lett.* **114**, 141801 (2015).
- [12] H. Fukuda, K. Harigaya, M. Ibe, and T. T. Yanagida, *Phys. Rev. D* **92**, 015021 (2015).
- [13] S. Dimopoulos, A. Hook, J. Huang, and G. Marques-Tavares, *J. High Energy Phys.* **11** (2016) 052.
- [14] J. E. Kim, *Phys. Rev. D* **31**, 1733 (1985).
- [15] L. Randall, *Phys. Lett. B* **284**, 77 (1992).
- [16] K.-w. Choi, *Phys. Rev. Lett.* **92**, 101602 (2004).
- [17] P. Svrcek and E. Witten, *J. High Energy Phys.* **06** (2006) 051.
- [18] B. Lillard and T. M. P. Tait, *J. High Energy Phys.* **11** (2017) 005.
- [19] B. Lillard and T. M. P. Tait, *J. High Energy Phys.* **11** (2018) 199.
- [20] P. Cox, T. Gherghetta, and M. D. Nguyen, *J. High Energy Phys.* **01** (2020) 188.
- [21] S. Dimopoulos, *Phys. Lett.* **84B**, 435 (1979).
- [22] B. Holdom and M. E. Peskin, *Nucl. Phys.* **B208**, 397 (1982).
- [23] M. Dine and N. Seiberg, *Nucl. Phys.* **B273**, 109 (1986).
- [24] J. M. Flynn and L. Randall, *Nucl. Phys.* **B293**, 731 (1987).
- [25] P. Agrawal and K. Howe, *J. High Energy Phys.* **12** (2018) 029.
- [26] P. Agrawal and K. Howe, *J. High Energy Phys.* **12** (2018) 035.
- [27] T. Gherghetta, N. Nagata, and M. Shifman, *Phys. Rev. D* **93**, 115010 (2016).
- [28] M. K. Gaillard, M. B. Gavela, R. Houtz, P. Quilez, and R. del Rey, *Eur. Phys. J. C* **78**, 972 (2018).
- [29] P. W. Graham, I. G. Irastorza, S. K. Lamoreaux, A. Lindner, and K. A. van Bibber, *Annu. Rev. Nucl. Part. Sci.* **65**, 485 (2015).
- [30] M. Tanabashi *et al.* (Particle Data Group), *Phys. Rev. D* **98**, 030001 (2018).
- [31] J. Jaeckel and M. Spannowsky, *Phys. Lett. B* **753**, 482 (2016).
- [32] S. Knapen, T. Lin, H. K. Lou, and T. Melia, *Phys. Rev. Lett.* **118**, 171801 (2017).
- [33] E. Izaguirre, T. Lin, and B. Shuve, *Phys. Rev. Lett.* **118**, 111802 (2017).
- [34] W. J. Marciano, A. Masiero, P. Paradisi, and M. Passera, *Phys. Rev. D* **94**, 115033 (2016).
- [35] I. Brivio, M. B. Gavela, L. Merlo, K. Mimasu, J. M. No, R. del Rey, and V. Sanz, *Eur. Phys. J. C* **77**, 572 (2017).
- [36] M. Bauer, M. Neubert, and A. Thamm, *J. High Energy Phys.* **12** (2017) 044.
- [37] B. Döbrich, F. Ertas, F. Kahlhoefer, and T. Spadaro, *Phys. Lett. B* **790**, 537 (2019).
- [38] N. Craig, A. Hook, and S. Kasko, *J. High Energy Phys.* **09** (2018) 028.
- [39] M. B. Gavela, R. Houtz, P. Quilez, R. del Rey, and O. Sumensari, *Eur. Phys. J. C* **79**, 369 (2019).
- [40] L. Merlo, F. Pobbe, S. Rigolin, and O. Sumensari, *J. High Energy Phys.* **06** (2019) 091.
- [41] M. Bauer, M. Neubert, S. Renner, M. Schnubel, and A. Thamm, *arXiv:1908.00008*.
- [42] W. Altmannshofer, S. Gori, and D. J. Robinson, *Phys. Rev. D* **101**, 075002 (2020).
- [43] J. Jaeckel, M. Jankowiak, and M. Spannowsky, *Phys. Dark Universe* **2**, 111 (2013).
- [44] A. Mariotti, D. Redigolo, F. Sala, and K. Tobioka, *Phys. Lett. B* **783**, 13 (2018).
- [45] X. C. Vidal, A. Mariotti, D. Redigolo, F. Sala, and K. Tobioka, *J. High Energy Phys.* **01** (2019) 113.
- [46] J. Beacham *et al.*, *J. Phys. G* **47**, 010501 (2020).
- [47] D. Aloni, Y. Soreq, and M. Williams, *Phys. Rev. Lett.* **123**, 031803 (2019).
- [48] G. Alonso-Álvarez, M. B. Gavela, and P. Quilez, *Eur. Phys. J. C* **79**, 223 (2019).
- [49] J. Ebadi, S. Khatibi, and M. M. Najafabadi, *Phys. Rev. D* **100**, 015016 (2019).
- [50] M. B. Gavela, J. M. No, V. Sanz, and J. F. de Trocóniz, *Phys. Rev. Lett.* **124**, 051802 (2020).
- [51] B. Doebrich, J. Jaeckel, F. Kahlhoefer, A. Ringwald, and K. Schmidt-Hoberg, *J. High Energy Phys.* **02** (2016) 018.
- [52] J. P. Chou, D. Curtin, and H. J. Lubatti, *Phys. Lett. B* **767**, 29 (2017).
- [53] M. J. Dolan, T. Ferber, C. Hearty, F. Kahlhoefer, and K. Schmidt-Hoberg, *J. High Energy Phys.* **12** (2017) 094.
- [54] V. V. Gligorov, S. Knapen, M. Papucci, and D. J. Robinson, *Phys. Rev. D* **97**, 015023 (2018).
- [55] A. Ariga *et al.* (FASER Collaboration), *Phys. Rev. D* **99**, 095011 (2019).
- [56] A. Berlin, S. Gori, P. Schuster, and N. Toro, *Phys. Rev. D* **98**, 035011 (2018).
- [57] J. L. Feng, I. Galon, F. Kling, and S. Trojanowski, *Phys. Rev. D* **98**, 055021 (2018).

- [58] D. Curtin *et al.*, *Rep. Prog. Phys.* **82**, 116201 (2019).
- [59] D. Aloni, C. Fanelli, Y. Soreq, and M. Williams, *Phys. Rev. Lett.* **123**, 071801 (2019).
- [60] B. Döbrich, J. Jaeckel, and T. Spadaro, *J. High Energy Phys.* **05** (2019) 213.
- [61] L. Harland-Lang, J. Jaeckel, and M. Spannowsky, *Phys. Lett. B* **793**, 281 (2019).
- [62] G. Aielli *et al.*, [arXiv:1911.00481](https://arxiv.org/abs/1911.00481).
- [63] C. Cornella, P. Paradisi, and O. Sumensari, *J. High Energy Phys.* **01** (2020) 158.
- [64] M. Bauer, M. Neubert, and A. Thamm, *Phys. Rev. Lett.* **119**, 031802 (2017).
- [65] G. Marques-Tavares and M. Teo, *J. High Energy Phys.* **05** (2018) 180.
- [66] C. A. Baker *et al.*, *Phys. Rev. Lett.* **97**, 131801 (2006).
- [67] P. Di Vecchia and G. Veneziano, *Nucl. Phys.* **B171**, 253 (1980).
- [68] R. Kallosh, A. Linde, D. Linde, and L. Susskind, *Phys. Rev. D* **52**, 912 (1995).
- [69] T. Banks and N. Seiberg, *Phys. Rev. D* **83**, 084019 (2011).
- [70] L. Del Debbio, L. Giusti, and C. Pica, *Phys. Rev. Lett.* **94**, 032003 (2005).
- [71] S. Aoki *et al.* (Flavour Lattice Averaging Group), *Eur. Phys. J. C* **80**, 113 (2020).
- [72] J. R. Ellis and M. K. Gaillard, *Nucl. Phys.* **B150**, 141 (1979).
- [73] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevLett.124.221801> for further discussions on theoretical aspects, collider phenomenology, and additional background considerations.
- [74] A. M. Uranga, *Nucl. Phys. B, Proc. Suppl.* **171**, 119 (2007).
- [75] N. Aghanim *et al.* (Planck Collaboration), [arXiv:1807.06209](https://arxiv.org/abs/1807.06209).
- [76] C. Kilic, T. Okui, and R. Sundrum, *J. High Energy Phys.* **02** (2010) 018.
- [77] If the fermions were in a larger representation of $SU(3)$, then their mass could be larger by Casimir factors.
- [78] M. Aaboud *et al.* (ATLAS Collaboration), *Eur. Phys. J. C* **78**, 250 (2018).
- [79] A. M. Sirunyan *et al.* (CMS Collaboration), *J. High Energy Phys.* **08** (2018) 130.
- [80] A. M. Sirunyan *et al.* (CMS Collaboration), *Phys. Rev. D* **98**, 112014 (2018).
- [81] G. Aad *et al.* (ATLAS Collaboration), *J. High Energy Phys.* **03** (2020) 145.
- [82] J. P. Lees *et al.* (BABAR Collaboration), *Phys. Rev. Lett.* **107**, 221803 (2011).
- [83] G. Aad *et al.* (ATLAS Collaboration), *J. High Energy Phys.* **01** (2013) 086.
- [84] M. Aaboud *et al.* (ATLAS Collaboration), *Phys. Rev. D* **95**, 112005 (2017).
- [85] S. Chatrchyan *et al.* (CMS Collaboration), *Eur. Phys. J. C* **74**, 3129 (2014).
- [86] S. Schael *et al.* (ALEPH, DELPHI, L3, OPAL, SLD, LEP Electroweak Working Group, SLD Electroweak Group, and SLD Heavy Flavour Group), *Phys. Rep.* **427**, 257 (2006).
- [87] J. de Blas, M. Ciuchini, E. Franco, S. Mishima, M. Pierini, L. Reina, and L. Silvestrini, *J. High Energy Phys.* **12** (2016) 135.
- [88] O. Adriani *et al.* (L3 Collaboration), *Phys. Lett. B* **292**, 472 (1992).
- [89] K. Mimasu and V. Sanz, *J. High Energy Phys.* **06** (2015) 173.
- [90] M. Aaboud *et al.* (ATLAS Collaboration), *Nat. Phys.* **13**, 852 (2017).
- [91] S. Knapen, T. Lin, H. K. Lou, and T. Melia, *CERN Proc.* **1**, 65 (2018).
- [92] A. Pierce, B. Shakya, Y. Tsai, and Y. Zhao, *Phys. Rev. D* **97**, 095033 (2018).
- [93] A. M. Sirunyan *et al.* (CMS Collaboration), *Phys. Lett. B* **797**, 134826 (2019).
- [94] G. Aad *et al.* (ATLAS Collaboration), *Eur. Phys. J. C* **76**, 210 (2016).
- [95] J. Alwall, M. Herquet, F. Maltoni, O. Mattelaer, and T. Stelzer, *J. High Energy Phys.* **06** (2011) 128.
- [96] T. Sjostrand, S. Mrenna, and P. Z. Skands, *J. High Energy Phys.* **05** (2006) 026.
- [97] T. Sjostrand, S. Mrenna, and P. Z. Skands, *Comput. Phys. Commun.* **178**, 852 (2008).
- [98] J. de Favereau, C. Delaere, P. Demin, A. Giammanco, V. Lemaitre, A. Mertens, and M. Selvaggi (DELPHES3 Collaboration), *J. High Energy Phys.* **02** (2014) 057.
- [99] Y. Gershtein, *Phys. Rev. D* **96**, 035027 (2017).
- [100] CMS Collaboration, First level track jet trigger for displaced jets at high luminosity LHC, Technical Report No. CMS-PAS-FTR-18-018, CERN, Geneva, 2018, <http://cds.cern.ch/record/2647987>.
- [101] M. Aaboud *et al.* (ATLAS Collaboration), *Phys. Rev. D* **97**, 052012 (2018).
- [102] A. M. Sirunyan *et al.* (CMS Collaboration), *J. Instrum.* **13**, P10034 (2018).
- [103] A. M. Sirunyan *et al.* (CMS Collaboration), *Phys. Rev. D* **99**, 032011 (2019).
- [104] J. Alimena *et al.*, [arXiv:1903.04497](https://arxiv.org/abs/1903.04497).
- [105] Y. Gershtein and S. Knapen, *Phys. Rev. D* **101**, 032003 (2020).
- [106] This track timing information can be further extrapolated back to the fitted vertex point, which is used to define the consistency of the tracks in the time direction (Cut6 and Cut7) in the next paragraph.
- [107] A. Hook, S. Kumar, Z. Liu, and R. Sundrum, ALPLHC, <https://gitlab.com/ZhenLiuPhys/alplhc>.