

Leading logs in QCD axion effective field theory

Gonzalo Alonso-Álvarez,^{a,b} Fatih Ertas,^c Joerg Jaeckel,^a Felix Kahlhoefer^c
and Lennert J. Thormaehlen^a

^a*Institut für theoretische Physik, Universität Heidelberg,
Heidelberg, Germany*

^b*McGill University Department of Physics & McGill Space Institute,
3600 Rue University, Montréal, QC, H3A 2T8, Canada*

^c*Institute for Theoretical Particle Physics and Cosmology (TTK), RWTH Aachen University,
D-52056 Aachen, Germany*

E-mail: galonso@physics.mcgill.ca, ertas@physik.rwth-aachen.de,
jjaeckel@thphys.uni-heidelberg.de, kahlhoefer@physik.rwth-aachen.de,
l.thormaehlen@thphys.uni-heidelberg.de

ABSTRACT: The axion is much lighter than all other degrees of freedom introduced by the Peccei-Quinn mechanism to solve the strong CP problem. It is therefore natural to use an effective field theory (EFT) to describe its interactions. Loop processes calculated in the EFT may however explicitly depend on the ultraviolet cutoff. In general, the UV cutoff is not uniquely defined, but the dimensionful couplings suggest to identify it with the Peccei-Quinn symmetry-breaking scale. An example are $K^+ \rightarrow \pi^+ + a$ decays that will soon be tested to improved precision in NA62 and KOTO and whose amplitude is dominated by the term logarithmically dependent on the cutoff. In this paper, we critically examine the adequacy of using such a naive EFT approach to study loop processes by comparing EFT calculations with ones performed in complete QCD axion models. In DFSZ models, for example, the cutoff is found to be set by additional Higgs degrees of freedom and to therefore be much closer to the electroweak scale than to the Peccei-Quinn scale. In fact, there are non-trivial requirements on axion models where the cutoff scale of loop processes is close to the Peccei-Quinn scale, such that the naive EFT result is reproduced. This suggests that the existence of a suitable UV embedding may impose restrictions on axion EFTs. We provide an explicit construction of a model with suitable fermion couplings and find promising prospects for NA62 and IAXO.

KEYWORDS: Beyond Standard Model, Effective Field Theories

ARXIV EPRINT: [2101.03173](https://arxiv.org/abs/2101.03173)

Contents

1	Introduction	1
2	Loop-induced rare decays in QCD axion EFT	4
3	Loop-induced rare decays in QCD axion models	8
3.1	DFSZ-type models	8
3.2	KSVZ-type models	9
4	An EFT-inspired QCD axion model	10
4.1	Low-energy effective model	10
4.2	UV completion of the model	12
4.3	Light- and heavy-quark mass matrix diagonalization	13
4.4	Low-energy effective couplings	17
5	Flavour and CP effects in the EFT-inspired QCD axion model	17
5.1	$K^+ \rightarrow \pi^+ + a$ decay rate	18
5.2	CP violation in axion interactions	20
6	Phenomenology of the model and discovery opportunities	21
6.1	Astrophysical limits and helioscope searches	21
6.2	Constraints and opportunities from $K^+ \rightarrow \pi^+ + a$	23
7	Conclusions and discussion	25
A	CP violation in pseudoscalar interactions	27
B	Variations of the model	28
C	Mass diagonalization and axion interactions	30
D	Counterterm contribution to the kaon decay width	38
E	EDM calculations	42

1 Introduction

Effective Field Theories (EFTs) are extremely useful tools to test new physics scenarios in a model-independent way. Their main ingredients are an often unknown cutoff energy scale and a set of low-energy fields and symmetries from which the operator expansion is constructed. Some examples of EFTs include the Weak Effective Theory [1–3], the Standard Model Effective Field Theory [4, 5], and the Higgs Effective Field Theory [6–8].

Nearly all phenomenological studies of axions and axion-like particles¹ are performed within an EFT where one considers the Standard Model (SM) extended by a pseudoscalar particle — the axion — that interacts with the SM via dimension-5 operators involving either two gauge bosons or two SM fermions.

The existence of a light pseudoscalar boson with the aforementioned properties can have potentially observable consequences in a wide variety of setups. Light axions may be copiously produced in astrophysical environments, allowing to search for them using helioscopes [9–11], cooling arguments [12, 13], or superradiance [14, 15]. Importantly, these light axions may constitute a fraction or the totality of the dark matter of the Universe [16–18], in which case a plethora of direct [9, 19–29] (see [30] for a recent overview) and indirect [31–45] detection experiments could be able to detect them. Axions at the heavier end² of the currently accessible mass range are best studied at the LHC [50–58] and B -factories [59–65]. For the intermediate regime between the MeV and the GeV scales, reactors [66], rare decay [59, 60, 63, 64, 67–71], beam dump/fixed target [72–80] experiments and long-lived particle detectors at the LHC [69, 81, 82] compete with supernovae [83–85] as the most favourable environments for probing axions (see [69] for a recent review). Of particular interest is the interplay between the various axionic couplings in different experiments and observations [85, 86] in these mass regimes — see [87–89] for the renormalization group (RG) equations that allow for a consistent comparison of constraints at different energy scales.

The rationale behind the employed EFTs is that the (QCD) axion is a pseudo-Goldstone boson arising from the Peccei-Quinn (PQ) solution to the Strong CP Problem [90–93]. It is expected that the underlying $U(1)_{\text{PQ}}$ symmetry is broken spontaneously at a very high energy scale while the axion remains very light. The pseudo-Goldstone boson nature of the axion then sets the guiding principles for the construction of its EFT. Firstly, the shift symmetry $a \rightarrow a + \text{const}$, resulting from the underlying $U(1)_{\text{PQ}}$, restricts the interactions of the axion to be of purely derivative form. Secondly, the $U(1)_{\text{PQ}}$ symmetry must be anomalous under QCD in order for the strong CP problem to be solved. Non-perturbative effects then generate an axion potential and in particular a non-vanishing mass. Moreover, the QCD anomaly and similar anomalies under the other gauge groups result in characteristic couplings of the axion to two gauge bosons. Finally, it seems natural to identify the cutoff scale of the EFT with the PQ symmetry-breaking scale, which is closely related to the axion decay constant, commonly denoted by f_a . This choice is suggested also by this scale appearing as the mass scale in the couplings of the dimension-5 derivative and two gauge boson interactions.

The adequacy of the EFT language to study axion phenomenology is based on the assumption that the presence of the axion is the only manifestation of the PQ construction at energies below the scale f_a . At tree level, it is usually sufficient that any other particles

¹Henceforth, we focus on the concrete case of (QCD) axions, though more general axion-like particles are usually motivated by similar symmetry considerations. Both particles are described by a generalized EFT that is based on the same construction principles, but for axion-like particles the requirement of addressing the strong CP problem of QCD is not enforced.

²See [46–49] for some recent work motivating the possibility that QCD axions could be heavier than expected.

are sufficiently heavy so that they cannot be produced at the energy scale relevant to the process under consideration. However, higher energy scales can become relevant at loop level where the contribution from off-shell particles in the loop needs to be included. In particular, when the EFT calculation yields divergent contributions, it may be more appropriate to identify the EFT cutoff scale with the mass of the lightest particle heavier than the axion instead of with f_a .

In order to construct a suitable axion EFT it is therefore essential to check whether the underlying model features additional relevant degrees of freedom below f_a . The goal of this paper is to do this in a number of full-fledged QCD axion models, thereby assessing the validity of studies performed using an EFT framework. As a particularly pertinent example, we focus on loop-induced flavour-violating decays involving the axion. Indeed, some of the one-loop diagrams involved in these processes are logarithmically divergent and therefore sensitive to the physics around the cutoff of the EFT. We show that the *leading log* prescription, where one focuses on the logarithmically divergent term and identifies the cutoff scale with f_a , commonly employed in EFT calculations, produces results which can differ qualitatively and quantitatively from the ones obtained using full models. The reason for this discrepancy, as will become clear, is that most popular (DFSZ- and KSVZ-type) QCD axion models [94–97] do not satisfy the assumptions that implicitly enter the EFT-based loop calculation. Our results are in line with the earlier work in a similar direction presented in [87]. There, a leading-order RG evolution was employed to show that simple electroweak-complete axion models (e.g. DSFZ-type models explicitly including the two Higgs doublets) feature an absence of a large logarithm in the flavour-changing effects. Our loop-based calculations are in agreement with these results.

In this situation, the question arises whether there exists an explicit field theoretic³ UV-complete QCD axion model that allows for large logarithms in flavour-changing observables. A positive answer amounts to finding ways to generate tree-level couplings between the axion and SM fermions without introducing new states below the scale of PQ symmetry breaking. We address this in two steps. First, we consider an effective model valid up to a scale $\gtrsim f_a$, which coincides with a specific charge assignment of the Lagrangian considered in ref. [87], where it was however not further explored. Our effective model also features some similarities with the flavoured axion models of [98, 99]. Crucially however, the charge assignments in our case are not flavour-dependent and thus tree-level flavour-violating interactions are not present in our setup. In the next step and drawing inspiration from Froggatt-Nielsen models [100], we construct a UV completion where the logarithms are explicitly calculable. The result is a QCD axion model that features flavour-conserving SM fermion couplings without the need for an extended Higgs sector, which therefore provides an affirmative answer to the question posed above.

The logarithmic enhancement of the $K^+ \rightarrow \pi^+ + a$ decay rate in this model makes it a particularly good candidate to be tested at experiments like NA62 [101] and KOTO [102] as well as future experiments such as KLEVER [103].⁴ Furthermore, the potential existence

³The authors of [87] considered string-motivated UV models for which flavour effects are significantly suppressed.

⁴For our concrete numerical examples we focus on charged kaon decays as measured by NA62.

of an axion-electron coupling increases the production of axions in the solar interior. This enhancement results in a larger expected axion flux in helioscopes like IAXO [11], which are projected to have sensitivity to the model in the $\mathcal{O}(10 \text{ meV})$ mass range.

In essence, the goal of the present study is to assess whether a given combination of axion EFT and cutoff is consistent with an embedding into a more fundamental theory. In the context of quantum gravity and string theory, the terms *landscape* and *swampland* have become popular to respectively denote the set of low energy effective theories that can and cannot consistently be incorporated into a theory of gravity [104]. Although our considerations are purely field-theoretic (see [105] for a swampland conjecture that is independent of the Planck scale), the landscape/swampland analogy serves as a motivation to raise the question of whether there are additional constraints on embeddable axion EFTs that have not yet been fully appreciated. In the case at hand, we find a promising avenue to construct the desired UV completion. That said, our model hints at some potentially non-trivial requisites. For example, achieving the absence of tree-level flavour-changing interactions in the EFT requires a non-trivial choice of the model parameters and potential tuning. Furthermore, realizing a sufficiently large top Yukawa coupling may constrain the separation of scales between the UV cutoff and f_a , as was noted in [87]. This suggests that mild additional assumptions (e.g. the absence of fine-tuning) may give strong constraints on the possible embeddings of axion EFTs into full UV models.

This paper is structured as follows. In section 2, we set up the QCD axion EFT under examination and calculate the expected rate for the $K^+ \rightarrow \pi^+ + a$ decay within this framework. This result is compared with detailed calculations in DFSZ- and KSVZ-type QCD axion models in section 3. After establishing the discrepancy between the EFT prescription and existing models, in section 4 we present a new QCD axion model that satisfies the assumptions under which the EFT is constructed. The phenomenology of this model is explored in sections 5 and 6, after which we conclude in section 7. A discussion of CP violation in axion couplings and most of the technical details can be found in the appendices A–E.

2 Loop-induced rare decays in QCD axion EFT

The most general effective Lagrangian describing all possible interactions of an axion (or axion-like particle) with the Standard Model fields before electroweak (EW) symmetry breaking and involving dimension-5 operators can be written in a compact form as [54, 106]

$$\mathcal{L} = -\frac{a}{f_a} \sum_F c_{FF} \frac{\alpha_F}{8\pi} F_{\mu\nu}^i \tilde{F}^{\mu\nu,i} + \frac{\partial_\mu a}{f_a} \sum_\chi \bar{\chi} C_\chi \gamma^\mu \chi, \quad (2.1)$$

where the first sum runs over SM gauge bosons with couplings c_{FF} . In the second sum, the SM chiral fermion multiplets are summarized in $\chi = (Q_L, L_L, u_R, d_R, e_R)$ and each of the C_χ is a matrix in generation space that allows for flavour changing effects (see appendix A for a discussion of the CP properties of C_χ). It is also customary to define the gauge boson couplings $g_{aFF} = c_{FF} \alpha_F / (2\pi f_a)$, but for our purposes the dimensionless couplings c_{FF} are more convenient. The dimensionful quantity f_a , usually known as the axion decay constant,

serves as the large scale for the expansion of the operators in the EFT. In eq. (2.1), there is an ambiguity in the definition of f_a and the EFT coefficients. As long as the coupling to gluons c_{gg} is nonzero, as is necessarily the case for any QCD axion, we can solve this ambiguity by normalizing f_a in such a way that $c_{gg} \equiv 1$, and we do so in the rest of this work.

The other relevant feature of the Lagrangian is that a potential for the axion is generated. For a nonzero coupling to gluons, non-perturbative QCD dynamics provide for a mass term [92, 95, 107, 108]

$$m_a \simeq \sqrt{\frac{m_\pi^2 f_\pi^2}{f_a^2} \frac{m_u m_d}{(m_u + m_d)^2}} \simeq 5.7 \mu\text{eV} \left(\frac{10^{12} \text{ GeV}}{f_a} \right). \quad (2.2)$$

At energies below the electroweak scale, it is convenient to rewrite the axion-fermion couplings in eq. (2.1) using the mass eigenbasis for the SM fermions,

$$\mathcal{L} \supset \frac{\partial_\mu a}{f_a} \left(\sum_{f=u,d,\ell,\nu} \bar{f}_L c_{f,L} \gamma^\mu f_L + \sum_{f=u,d,\ell} \bar{f}_R c_{f,R} \gamma^\mu f_R \right), \quad (2.3)$$

where the sum runs over the up, down, charged lepton, and neutrino flavour-triplets u , d , ℓ , and ν . Each f thus summarizes the three generations. The coupling matrices $c_{f,L}$ and $c_{f,R}$ for left- and right-handed fields are related to the initial C_χ matrices in eq. (2.1) by unitary matrices [54]. In the left-handed quark sector, we have the additional relation $c_{d,L} = V^\dagger c_{u,L} V$ involving the CKM matrix V , as both up- and down-type quarks are part of the same left-handed doublet. As neutrinos are massless in the SM, we can choose the unitary transformation of left-handed neutrino fields to equal the one of their charged-lepton counterparts so that $c_{\ell,L} = c_{\nu,L}$. Note that we could have also written down the interactions in a vector/axial-vector basis with the relevant matrices reading $c_{f,V} = (c_{f,R} + c_{f,L})/2$ and $c_{f,A} = (c_{f,R} - c_{f,L})/2$, respectively.

The Lagrangian as written down in eq. (2.3) in principle allows for arbitrary flavour-coupling structures. While some axion models [98, 99] feature flavour-dependent and even flavour non-diagonal interactions (see [71] for a summary of the associated phenomenology), they are not the main focus of this work. That said, as such couplings may also arise in our explicit model construction, we discuss them in some detail later.

For now, let us consider flavour-universal couplings in the quark and lepton sectors. This is already sufficient to demonstrate our main point, namely that the leading logarithm in FCNCs obtained in the naive EFT does, in many cases, not agree with the result of more complete models. In this case, we can remove the vectorial part of the quark and lepton couplings by performing universal vectorial phase rotations of the quark and lepton fields. These leave the mass terms and the charged current interactions with W^\pm bosons invariant and therefore no scalar Yukawa couplings to a appear. The only possible effects are two-gauge-boson couplings arising for the chirally coupled electroweak fields, as noted in [71], corresponding to the anomaly for the baryon and lepton number U(1) symmetries. The effect can therefore be absorbed in a redefinition of the corresponding two-gauge-boson coupling coefficient, most notably c_{WW} . As we discuss below, the contribution from c_{WW} to the loop processes studied in this work is finite and therefore a shift in this coefficient




Figure 1. One-loop diagrams inducing the flavour-violating $s \rightarrow d + a$ transition. We include the self-energy contribution on the external strange-quark leg (analogously for the down-quark line) arising from the renormalization of quark fields, see appendix D. All Feynman diagrams throughout this work are drawn using the TikZ-Feynman [111] package.

does not play a role when evaluating the leading logarithm. We thus only need to consider axial-vector couplings of the axion to quarks and leptons,⁵

$$\mathcal{L} \supset \frac{\partial_\mu a}{2f_a} c_q \sum_{f=u,d} \bar{f} \gamma^\mu \gamma_5 f + \frac{\partial_\mu a}{2f_a} c_l \sum_{f=\ell,\nu} \bar{f} \gamma^\mu \gamma_5 f, \quad (2.4)$$

where the neutrino fields are understood to be purely left handed, $\nu = \nu_L$. With these assumptions, no inter-generation fermion couplings and therefore no flavour-changing processes are present at tree level. But even if these are absent, flavour-changing neutral currents (FCNCs) appear at one loop through the diagrams depicted in figure 1. Rare decays induced by such transitions have been exploited in the literature to test the axion paradigm [59, 60, 63, 64, 69–71], and lepton-flavour violation has been studied in similar detail [109, 110].

In this work, we perform a detailed study of the FCNC transition in figure 1, with a focus on the impact of the possible UV completions of the axion effective Lagrangian in eq. (2.4). For concreteness and also because it is of significant experimental relevance, we concentrate on the kaon decay $K^+ \rightarrow \pi^+ + a$. That said, our results are also applicable to other FCNC-induced processes.

The $s \rightarrow d + a$ transition can be described using the effective Hamiltonian

$$\mathcal{H}_{s \rightarrow da} = \partial_\mu a \bar{d} \gamma^\mu (h_{ds}^S + h_{ds}^P \gamma_5) s + \text{h.c.} \quad (2.5)$$

The hadronic matrix element for the $K^+ \rightarrow \pi^+ + a$ decay⁶ can then be parameterized as [112]

$$\langle \pi(p') | \bar{d} \gamma^\mu (h_{ds}^S + h_{ds}^P \gamma_5) s | K(p) \rangle = h_{ds}^S P^\mu f_+(q^2) + h_{ds}^P q^\mu f_-(q^2), \quad (2.6)$$

where $P = p + p'$ and $q = p - p'$ so that the momentum transfer is $q^2 = m_a^2$. For a sufficiently light axion, we only need the first form factor at $q^2 \simeq 0$. Recent Lattice QCD evaluations [113] give $f_+(0) = 0.9706(27)$. In this limit, the decay width can be expressed as

$$\Gamma(K^+ \rightarrow \pi^+ a) = \frac{|h_{ds}^S|^2}{16\pi m_{K^+}^3} (m_{K^+}^2 - m_{\pi^+}^2)^2 \lambda^{1/2}(m_{K^+}^2, m_{\pi^+}^2, m_a^2) f_+^2(m_a^2), \quad (2.7)$$

where we have introduced the notation $\lambda(x, y, z) = x^2 + y^2 + z^2 - 2(xy + xz + yz)$.

⁵In DFSZ models, the couplings for up- and down-type quarks can actually differ even in absence of an explicit flavour structure due the existence of an extended EW sector and mixing effects. In section 3.1 we avoid this complication by adapting the calculation of [59] which is done in the Yukawa basis.

⁶As already mentioned, we focus on the charged kaon decay as studied in NA62 [101], but a similar calculation can be done for neutral kaon decays (KOTO [102] and KLEVER [103]).

Within the framework of eq. (2.4), the coefficients $h_{ds}^{S,P}$ can be expressed in terms of the parameters of the axion EFT. Given that, as we have seen, the QCD matrix element for h_{ds}^P vanishes due to parity, we only need [60, 114]

$$\begin{aligned}
 h_{ds}^S = & -\frac{G_F}{16\sqrt{2}\pi^2} \frac{1}{f_a} c_q \sum_{q=u,c,t} V_{qd}^* V_{qs} m_q^2 \left(\log\left(\frac{\Lambda^2}{m_q^2}\right) - \frac{m_q^4 - 8m_q^2 m_W^2 + 7m_W^4 + 6m_W^4 \log\left(\frac{m_q^2}{m_W^2}\right)}{2(m_q^2 - m_W^2)^2} \right) \\
 & - \frac{3G_F^2 m_W^4}{\pi^2} \frac{c_{WW}}{32\pi^2 f_a} \sum_{q=u,c,t} V_{qd}^* V_{qs} f(m_q^2/m_W^2),
 \end{aligned} \tag{2.8}$$

where we have substituted the divergence arising from the calculation of the diagrams in figure 1 by a leading logarithm depending on an a priori undetermined UV scale Λ . In an EFT sense, this identification can also be understood by considering a leading-order RG evolution between the high and the low scale [87, 89]. We have also introduced the loop function [60]

$$f(x) = \frac{x(1 + x(\ln(x) - 1))}{(1 - x)^2}, \tag{2.9}$$

and expanded the amplitude at leading order in the external momenta (corresponding to dropping terms suppressed by m_d/m_W , m_s/m_W , and m_a/m_W) in order to determine the leading finite contribution. This expansion, however, has barely any effects compared to the “leading log”. Finally, we have neglected the contribution from the gluon anomalous coupling, which can be computed based on mixing effects [47, 115] but is subleading to the leading logarithmic term.

The result above has been computed in unitary gauge using the basis of eq. (2.4), where the QCD axion couples only derivatively. An axion-dependent rotation of the fermion fields allows to trade the derivative terms for pseudoscalar ones $\partial_\mu a \bar{f} \gamma^\mu \gamma_5 f \rightarrow 2m_f a \bar{f} i \gamma_5 f$, importantly along with anomaly terms. One can check that the pseudoscalar interactions give the same leading log result as the axial-vector ones. In Feynman gauge, a factor of 4 difference between the axial-vector and pseudoscalar interaction has been reported (see footnote 3 in ref. [67]), but we have confirmed that the claimed discrepancy disappears when the EW Goldstone bosons are taken into account.

Combining the above equations, one can evaluate the decay rate as a function of the EFT parameters and, importantly, of the UV cutoff Λ . This can in principle be used to place bounds on f_a using the experimental constraint $\text{Br}(K^+ \rightarrow \pi^+ a) < 7.3 \cdot 10^{-11}$ from the E787 and E949 experiments at BNL [116], which NA62 expects to improve by an order of magnitude by 2025 [117, 118]. In practice, the reach of the constraint depends significantly on the choice of the value of Λ . It is common to identify Λ with the only other scale that is available in the EFT description, which is the decay constant f_a [63, 67, 114]. However, this is an *ad hoc* choice which, as we will see, is not reproduced in the most popular UV completions of the axion EFT. In order to clarify this issue, we now turn to study the $s \rightarrow d + a$ transition in the benchmark KSVZ and DFSZ QCD axion models.

3 Loop-induced rare decays in QCD axion models

3.1 DFSZ-type models

Let us compute the flavour-violating decay of the kaon through the diagrams of figure 1 in DFSZ [96, 97] models. Crucially, we use a complete DFSZ model explicitly including the two Higgs doublets in order to regularize the logarithmic divergence and identify the cutoff scale Λ . The results are obtained from [59], where the $B \rightarrow K + a$ process was studied, by applying the obvious substitutions. Before going into the details, let us note that these calculations are performed in the basis where the axion interactions are of Yukawa and not of derivative type. As we briefly discuss below eq. (3.8), this avoids complications arising when separate chiral rotations of up- and down-type quarks are performed.

Our starting point is the simplest DFSZ model, which is a two-Higgs-doublet model (2HDM) of Type II extended by a singlet complex scalar Φ that transmits its PQ charge via the operator

$$\mathcal{L} \supset \Phi^2 H_u H_d. \tag{3.1}$$

The PQ symmetry is broken when Φ acquires a vacuum expectation value (VEV) f_a , which we assume to be much larger than the EW VEV $v = (v_u^2 + v_d^2)^{1/2}$. As usual, we denote $\tan \beta = v_u/v_d$. The axion inherits couplings to the SM quarks due to its mixing with the pseudoscalar Higgs, resulting in axion Yukawa interactions with the up $\sim (m_{u,i}/f_a) \cos^2 \beta/3$ and down $\sim (m_{d,i}/f_a) \sin^2 \beta/3$ quarks. With this, the value of h_{ds}^S in the effective Hamiltonian eq. (2.5) describing the $s \rightarrow d + a$ transition is found to be

$$h_{ds}^S = -\frac{G_F}{16\pi^2} \frac{\cos^2 \beta}{3f_a} \sum_{q=u,c,t} V_{qd}^* V_{qs} m_q^2 (X_1^q + X_2^q \cot^2 \beta). \tag{3.2}$$

The sum runs over all the up-type quark flavours and the (finite) loop functions are given by [59]

$$X_1^q = 2 + \frac{m_{H^\pm}^2}{m_{H^\pm}^2 - m_q^2} - \frac{3m_W^2}{m_q^2 - m_W^2} + \frac{3m_W^4 (m_{H^\pm}^2 + m_W^2 - 2m_q^2)}{(m_{H^\pm}^2 - m_W^2) (m_q^2 - m_W^2)^2} \log \frac{m_q^2}{m_W^2} \\ + \frac{m_{H^\pm}^2}{m_{H^\pm}^2 - m_q^2} \left(\frac{m_{H^\pm}^2}{m_{H^\pm}^2 - m_q^2} - \frac{6m_W^2}{m_{H^\pm}^2 - m_W^2} \right) \log \frac{m_q^2}{m_{H^\pm}^2}, \tag{3.3}$$

$$X_2^q = -\frac{2m_q^2}{m_{H^\pm}^2 - m_q^2} \left(1 + \frac{m_{H^\pm}^2}{m_{H^\pm}^2 - m_q^2} \log \frac{m_q^2}{m_{H^\pm}^2} \right), \tag{3.4}$$

in terms of the charged Higgs boson mass m_{H^\pm} . The different sign of the leading logarithmic term in the limit of large m_{H^\pm} compared to eq. (2.8) is due to a negative c_q in the DFSZ model. We can then use eq. (2.7) to compute the expected $K^+ \rightarrow \pi^+ + a$ branching ratio in DFSZ models. The expected NA62 constraints [117] as a function of m_{H^\pm} and $\tan \beta$ are shown in figure 2. Note that for each value of $\tan \beta$, there is a value of $m_{H^\pm}^\pm$ for which h_{ds}^S in eq. (3.2) changes sign, which causes the funnel in figure 2 along which the bound disappears.

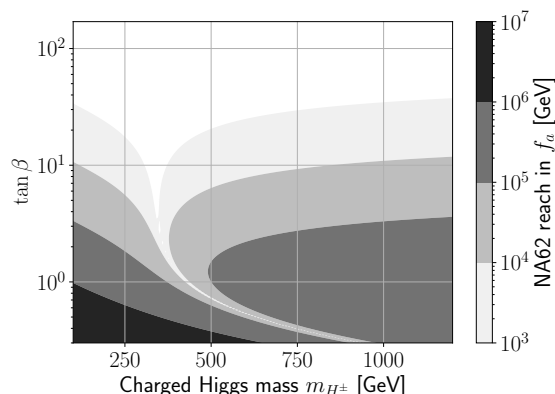


Figure 2. Expected NA62 reach on f_a from $K^+ \rightarrow \pi^+ + a$ in flavour-universal DFSZ models, in terms of the 2HDM parameters m_{H^\pm} and $\tan \beta$. We use the projection of [117] for $\text{Br}(K^+ \rightarrow \pi^+ + a)$.

With this computation, it becomes explicit that the value of the cutoff scale of the leading logarithm for the DFSZ model is the charged Higgs boson mass, $\Lambda_{\text{DFSZ}} = m_{H^\pm}$, which is orders of magnitude smaller than f_a in the interesting region of parameter space. Moreover, cancellations in the decay rate can occur for specific parameter points, an effect that is not captured within the EFT approach. Finally, FCNCs can also be suppressed by large values of $\tan \beta$ [87]. Overall, this means that the $K^+ \rightarrow \pi^+ + a$ branching ratio in DFSZ models is much smaller than expected based on the EFT computation. As a consequence, the actual bounds and projections from experiments such as NA62 on the DFSZ axion parameter become correspondingly weaker. This comparison can be seen in figure 5 in section 6, where the dark region represents the branching ratio for a benchmark value of $m_{H^\pm} = 800 \text{ GeV}$ and $1 \leq \tan \beta \leq 5$, roughly representative of the experimentally allowed 2HDM parameter space (see e.g. [119, 120]). We also display a lighter region which corresponds to only requiring perturbativity of the Yukawa couplings, $0.25 \leq \tan \beta \leq 170$ [121].

3.2 KSVZ-type models

The computation above shows that the EFT calculation with the naive leading-log approximation is not appropriate for DFSZ axion models. The reason for this is that DFSZ axion models introduce new degrees of freedom (the second Higgs doublet) at a scale below the naive EFT cutoff f_a . In KSVZ [94, 95] models, one expects that the EFT leading log term should not appear either. The reason is that in KSVZ models no tree-level fermion couplings are present, and even if they can be generated by an axion-dependent chiral rotation of fermion fields, the logarithmically divergent contribution from the 1-loop process involving the derivative coupling $\partial_\mu a / (2f_a) \bar{f} \gamma^\mu \gamma^5 f$ exactly cancels with the one coming from the pseudoscalar coupling $ia / f_a m_f \bar{f} \gamma^5 f$, so that only a finite piece remains. From a different perspective, this cancellation reflects the equivalence of the linear and polar representation used for the complex scalar field containing the axion [122].

To check this statement, we start with

$$\mathcal{L} \supset -\frac{a}{f_a} \frac{\alpha_s}{8\pi} G^{a,\mu\nu} \tilde{G}_{\mu\nu}^a, \quad (3.5)$$

and perform a chiral rotation [54] on *all* SM quarks, i.e. $q = (u, d, s, c, b, t)^T$, with

$$q \rightarrow \exp\left(i\kappa_q \frac{a}{2f_a} \gamma_5\right) q, \quad (3.6)$$

where κ_q is a 6×6 diagonal matrix. Taking, e.g., $\text{Tr}[\kappa_q] = 1$ removes the gluon coupling completely [54]. As the one-loop process for $s \rightarrow d + a$ is $\mathcal{O}(g^2)$, we have to collect all terms that contribute at the same order. The relevant terms before the rotation are

$$\mathcal{L} \supset \bar{q} i \gamma^\mu \partial_\mu q - \bar{q} M_q q - \frac{g}{\sqrt{2}} \bar{u}_L \gamma^\mu W_\mu^+ V d_L + \text{h.c.}, \quad (3.7)$$

yielding after the chiral rotation (expanded at LO in the ALP field)

$$\mathcal{L} \supset -\frac{\partial^\mu a}{2f_a} \bar{q} \kappa_q \gamma_\mu \gamma_5 q - \frac{ia}{f_a} \bar{q} \kappa_q M_q \gamma_5 q - \frac{ia g}{2\sqrt{2}f_a} \bar{u}_L \gamma^\mu W_\mu^+ (\kappa_u V - V \kappa_d) d_L + \text{h.c.}, \quad (3.8)$$

where M_q is the quark mass matrix and κ_u and κ_d are the 3×3 submatrices of κ_q including only the elements for the up-type and down-type quarks, respectively. After tracing the influence of these terms on $s \rightarrow d + a$ at one loop, one notices that the UV-divergent terms cancel exactly. Note that in this context the third term in eq. (3.8) provides crucial contributions by axion emissions off the $W - q - q'$ vertices, which do not vanish unless κ_q is fully flavour-universal.⁷ We can therefore robustly conclude that a leading log does not appear in KSVZ models at the one-loop level.

4 An EFT-inspired QCD axion model

4.1 Low-energy effective model

Because the naive EFT result is not recovered in common axion models like KSVZ or DFSZ, one wonders if there is any UV completion of the axion EFT where the leading-log result is actually applicable, meaning that the full result of the kaon decay width is also enhanced by as large a logarithm (or an even larger one) as in (2.8). From our discussion above, we can identify two key requirements for such a model. First, it should not contain any additional new states below the new-physics scale Λ apart from the axion (unlike the DFSZ model). And second, it should feature physical couplings to quarks that cannot be removed by a chiral rotation (unlike the KSVZ model).

A general class of effective QCD axion models that satisfy both of these requirements is given by the Lagrangian⁸

$$\mathcal{L} \supset -\frac{\Phi}{\Lambda_{ij}^u} \bar{Q}_{Li} \tilde{H} u_{Rj} - \frac{\Phi}{\Lambda_{ij}^d} \bar{Q}_{Li} H d_{Rj} + \text{h.c.}, \quad (4.1)$$

⁷Intriguingly, this means that for DFSZ-like models such terms can in general not be neglected upon performing a chiral rotation from the Yukawa to the derivative basis.

⁸As noted before, this model is a member of the class of effective Lagrangians considered in [87] based on electroweak symmetry considerations. Ref. [87] also provided a string-related axion-like particle realisation to give a large logarithm. However, in that setup the FCNC is parametrically suppressed by $1/(16\pi^2) \sim 10^{-2}$ compared to the field theoretic construction presented below.

where Q_L , u_R , d_R , and H denote the usual SM chiral fermion fields and Higgs doublet before EW symmetry breaking, and the indices i, j label the three generations.

The complex scalar $\Phi = \phi/\sqrt{2} e^{ia/\langle\phi\rangle}$ contains the axion as its angular degree of freedom, and we have therefore normalized its PQ charge to $\chi_\Phi = 1$. The Higgs doublet has no PQ charge⁹ and the charges of the quarks therefore satisfy $\chi_{Q_{L_i}} - \chi_{u_{R_i}} = \chi_{Q_{L_i}} - \chi_{d_{R_i}} = 1$.

In order to avoid having to impose any restrictions on the coupling matrices $\Lambda^{u,d}$, we assume that the PQ charges are flavour independent,

$$\chi_L = \chi_{Q_{L_i}} \quad \text{and} \quad \chi_R = \chi_{d_{R_i}} = \chi_{u_{R_i}} \quad \text{for all generations.} \quad (4.2)$$

Note that this is the essential difference to the models considered in [98, 99], where the charges were assumed to be flavour-dependent in order to explain the observed flavour structure of the Standard Model Yukawa couplings. Such a choice generically leads to tree-level flavour-violating interactions for the axion. As our focus is on the loop-induced flavour effects, we choose flavour-independent charges to avoid these complications.

The QCD anomaly induced by the SM quarks is

$$N = \sum_f (\chi_{f_L} - \chi_{f_R}) T(R_f) = 3, \quad (4.3)$$

as in the usual DFSZ models. The electromagnetic anomaly coefficient,

$$E = \sum_f (\chi_{f_L} - \chi_{f_R}) Q_f^2, \quad (4.4)$$

can either be 8 or 5 depending on whether the leptons are PQ-charged or not (leptons can also couple to Φ^* , resulting in a negative contribution to E). Once the PQ scalar acquires a VEV $\langle\Phi\rangle = \langle\phi\rangle/\sqrt{2}$, we can write

$$\mathcal{L} \supset -Y_{ij}^u e^{i\frac{a}{\langle\phi\rangle}} \bar{Q}_{L_i} \tilde{H} u_{R_j} - Y_{ij}^d e^{i\frac{a}{\langle\phi\rangle}} \bar{Q}_{L_i} H d_{R_j} + \text{h.c.}, \quad (4.5)$$

where the SM Yukawa couplings are $Y_{ij}^{u,d} = \langle\phi\rangle/(\sqrt{2}\Lambda_{ij}^{u,d})$. Because the axion is automatically aligned with the Yukawas, there are no tree-level flavour-violating axion couplings in this model.

After EW symmetry breaking and CKM diagonalization, the usual pseudoscalar couplings of the axion to quarks,

$$\mathcal{L} \supset -\sum_q m_q \bar{q} e^{i\frac{a}{\langle\phi\rangle}} \gamma^5 q, \quad (4.6)$$

are recovered, meaning that the couplings are strictly proportional to the fermion masses.

The presence of axion-quark couplings in this model implies the existence of a divergent contribution to flavour-changing amplitudes at one loop from diagrams as the ones shown

⁹The model does not change if the Higgs is allowed to carry charge. The reason is that the five charges of Q_L , u_R , d_R , H , and Φ are restricted by the two interaction terms in eq. (4.1) to three conserved global U(1) symmetries which can be identified as hypercharge, baryon number, and the PQ symmetry. In this three-dimensional space, a direction for the PQ symmetry not involving the Higgs field can always be found.

in figure 1. When replacing this divergence by a leading logarithm, the only possibilities are $\log(\langle\phi\rangle^2/m_q^2)$ or $\log(\Lambda_{ij}/m_q^2)$. Since all of these scales are at least as large as the PQ breaking scale, we therefore necessarily obtain the large logarithmic enhancement of flavour-changing processes anticipated from the axion EFT.

Since we started from a non-renormalizable Lagrangian in eq. (4.1), it is not possible to obtain an exact result for the leading logarithm and as such one also lacks its physical interpretation. It is therefore instructive to understand how the effective interactions discussed above can be embedded in a UV-complete model, which allows to calculate the flavour-changing processes in detail.

4.2 UV completion of the model

In order to embed the effective model introduced above into a renormalizable UV completion, we add three generations of heavy coloured up- and down-type fermions, F_i^u and F_i^d , which resemble the messenger fields in the Froggatt-Nielsen mechanism [100]. For simplicity, we first study the case of a single generation, and later generalize it to three generations.

For one generation, the relevant renormalizable Lagrangian involves the terms

$$\begin{aligned} \mathcal{L} \supset & -\alpha^u \bar{Q}_L \tilde{H} F_R^u - \beta^u \bar{F}_L^u \Phi u_R + \text{h.c.} \\ & -\alpha^d \bar{Q}_L H F_R^d - \beta^d \bar{F}_L^d \Phi d_R + \text{h.c.}, \end{aligned} \tag{4.7}$$

where $\alpha^{u,d}$ and $\beta^{u,d}$ denote the couplings in the up- and down-type sectors.¹⁰

More generally, one can allow the possibility to use either Φ or Φ^* . This leads to slightly different models, which are described in appendix B. We focus on the example of the model as specified in eq. (4.7), but the other options can be treated in a similar way. For all terms to be gauge invariant, F^u and F^d must transform under the SM gauge groups like u_R and d_R , respectively, and left- and right-handed F fields (we omit the u/d superscript when referring to any of the two fields) must transform identically. The terms in eq. (4.7) always enforce the conditions $\chi_{Q_L} - \chi_{F_R} = 0$ and $\chi_{F_L} - \chi_{q_R} = 1$. It is precisely the sum of these equations that appears in the anomaly coefficients N and E in eqs. (4.3) and (4.4), as up- and down-type F and SM quarks couple similarly to gluons and photons. Hence, the additional SU(3) quarks do not change the anomaly coefficients $N = 3$ and $E = 8$ or 5 . Furthermore, a bare mass term for the F quarks,

$$\mathcal{L} \supset -\lambda \frac{\langle \Sigma \rangle}{\sqrt{2}} \bar{F}_L F_R + \text{h.c.}, \tag{4.8}$$

can be added. For easier extension to the multi-generation case, we have parameterized the mass in terms of the VEV of a real spurion field Σ for which $\chi_\Sigma = 0$. This imposes the additional condition $\chi_{F_L} - \chi_{F_R} = 0$. With this, all axial charges of all combinations of SM

¹⁰The only other dimension-4 term that is invariant under all symmetries considered is $\Phi\Phi^*HH^\dagger$. If included, such a term would generate a large contribution to the Higgs mass unless it comes with an extremely small prefactor. This is nothing but the usual naturalness problem that is common to all axion models, and which we do not seek to address in this work.

and F quarks have been fixed and the only remaining freedom is a shift of all charges by an arbitrary constant.

It is clear that after PQ and EW symmetry breaking, the terms in eq. (4.7) induce mass mixing between the different fields involved. In unitary gauge, we can write the mixing in matrix form as

$$\mathcal{L} \supset - \left(\bar{q}_L \ \bar{F}_L \right) M \begin{pmatrix} q_R \\ F_R \end{pmatrix} + \text{h.c.}, \quad (4.9)$$

where M depends on the VEVs of all scalar fields in our theory (H, Φ, Σ), the corresponding coupling matrices (α, β, λ), and the axion field a .

4.3 Light- and heavy-quark mass matrix diagonalization

In order to recover the quark masses and couplings, we proceed exactly as in the Standard Model and diagonalize the mass matrices M^u and M^d by unitary transformations of left- and right-handed fields. Only the main steps and results are mentioned here: we refer to appendix C for detailed calculations. Throughout this and the next subsection, we drop the labels u/d whenever the procedure is identical for both types of quarks. We consider at every point all three generations of SM and F quarks even though the generation indices are also omitted.

The full mass matrix M is given by

$$M = \frac{\langle \Sigma \rangle}{\sqrt{2}} \begin{pmatrix} 0 & \epsilon \epsilon' \alpha \\ \epsilon' \beta e^{ia/\langle \phi \rangle} & \lambda \end{pmatrix}, \quad (4.10)$$

where we have introduced two parameters

$$\epsilon = \frac{v}{\langle \phi \rangle} \quad \text{and} \quad \epsilon' = \frac{\langle \phi \rangle}{\langle \Sigma \rangle}, \quad (4.11)$$

with the Higgs VEV $v = 246$ GeV, which highlight the hierarchies of scales in our model. We can safely assume that the parameter ϵ is much smaller than one. In contrast, ϵ' may but does not need to be smaller than one for a general axion model, where extra degrees of freedom other than Φ may appear close to the PQ scale. Importantly, such additional degrees of freedom close to the PQ scale can lead to sizable corrections to our desired effective model in eq. (4.5), introducing e.g. tree-level flavour-violating axion couplings. In order to avoid this complication, we assume that the additional degrees of freedom at hand (the F quarks) are heavy enough so that these effects are sufficiently suppressed. In practice, this means that we require

$$\epsilon'' = \frac{\langle \phi \rangle}{\min m_{F_i}} \ll 1, \quad (4.12)$$

where $\min m_{F_i}$ is the mass of the lightest fermion that is not part of the SM. More details, such as quantifying how small ϵ'' has to be in the most general case as well as how tree-level flavour violation can be avoided even for sizeable ϵ'' , can be found in appendix C.

Continuing with the diagonalization, we define U to be a unitary matrix which diagonalizes the hermitian product MM^\dagger ,

$$U^\dagger MM^\dagger U = \Lambda^2 = \begin{pmatrix} M_q^2 & 0 \\ 0 & M_F^2 \end{pmatrix}. \quad (4.13)$$

Here, M_q and M_F are the diagonal mass matrices of Q and F quarks. Transforming the left-handed quark fields with U and the right-handed ones with $S = M^\dagger U \Lambda^{-1}$,

$$\begin{pmatrix} q_L \\ F_L \end{pmatrix} \rightarrow U \begin{pmatrix} q_L \\ F_L \end{pmatrix}, \quad \begin{pmatrix} q_R \\ F_R \end{pmatrix} \rightarrow S \begin{pmatrix} q_R \\ F_R \end{pmatrix}, \quad (4.14)$$

exactly gives us a basis of fields in which the mass matrix is diagonal and independent of a . Note that to lighten the notation we are using identical symbols for fields in the original Lagrangian, eq. (4.7), and for the mass eigenstates after the diagonalization.

By expanding U in ϵ and ϵ'' , we can express the quark masses in terms of the UV parameters. At leading order, the masses are

$$M_q^2 = \frac{v^2}{2} \epsilon'^2 U_\delta^\dagger (\alpha \lambda^{-1} \beta \beta^\dagger \lambda^{\dagger-1} \alpha^\dagger) U_\delta, \quad (4.15)$$

$$M_F^2 = \frac{\langle \Sigma \rangle^2}{2} U_\xi^\dagger (\lambda \lambda^\dagger) U_\xi, \quad (4.16)$$

where U_δ and U_ξ are unitary matrices which diagonalize the hermitian matrices in parentheses. It is important to keep in mind that all matrices in (4.15) and (4.16) exist for up- and down-type particles, giving a total of 12 different quark masses. Any physical realization of the coupling matrices needs to reproduce the masses of SM quarks in eq. (4.15). Since the mass of the top quark is close to $\langle H \rangle = v/\sqrt{2}$, it may be difficult or impossible to recover such a high value when all F quarks are much heavier than the PQ scale and when simultaneously requiring perturbativity of all Yukawa couplings. A similar observation was made in [87], where it was concluded that the cutoff of a Lagrangian as the one in (4.1) can at most be of order f_a . However, in our explicit UV completion the masses of the six F quarks can have a large intrinsic hierarchy, which can in principle lead to logarithmic enhancements of flavour-violating effects by factors larger than $\log(f_a^2/m_q^2)$ and possibly also tree-level flavour violation, and simultaneously reproduce the observed top Yukawa. These effects are part of the non-trivial flavour properties of our model, whose detailed investigation goes beyond the scope of this work. In the main part of this text, we work in the limit $\epsilon'' \ll 1$ while remaining agnostic about the exact flavour structure that may realize this hierarchy. As a proof of principle, in appendix C we give an explicit example with $\epsilon'' = 0.2$ for which all of our calculations apply and the SM is reproduced.

Axion couplings to gauge bosons. The matrices U and S introduced in the previous section generally contain axion-dependent phases. As a consequence, the quark-field transformations in eq. (4.14), which are required to diagonalize the mass matrix, source axion interaction terms. To understand these better, it is helpful to split the diagonalization procedure into two subsequent field redefinitions, of which only the first one depends on the axion field.

Looking at the full UV Lagrangian in (4.7), it is clear that the a dependence of the mass-mixing terms can be absorbed into the quark fields by the redefinitions

$$u_R \rightarrow e^{-\frac{ia}{\langle\phi\rangle}} u_R, \quad d_R \rightarrow e^{-\frac{ia}{\langle\phi\rangle}} d_R. \quad (4.17)$$

This transformation removes the axion field a from M in eq. (4.10), while the quark kinetic terms generate derivative couplings of the axion to right-handed quarks. Hence, both diagonalization matrices U and S can subsequently be chosen to be independent of the axion field. The resulting derivative interactions are explicitly written out below in terms of the mass-diagonal quark fields (see eq. (4.22)).

In addition to the kinetic terms, the path integral measure is not invariant under the transformation in (4.17) and anomalous interaction terms between the axion and gauge bosons arise. The $U(1)_A \times G \times G$ anomaly, with G being either the strong or the EW gauge group, sources the interaction terms [90, 91, 123–125]

$$\begin{aligned} \mathcal{L} \supset & -2N \cdot \frac{\alpha_s}{16\pi \langle\phi\rangle} a \epsilon^{\mu\nu\alpha\beta} G_{\mu\nu}^a G_{\alpha\beta}^a - E \cdot \frac{\alpha_{\text{EM}}}{8\pi \langle\phi\rangle} a \epsilon^{\mu\nu\alpha\beta} F_{\mu\nu} F_{\alpha\beta} \\ & + E \cdot \frac{\alpha_{\text{EM}} s_W}{4\pi c_W} \frac{a}{\langle\phi\rangle} \epsilon^{\mu\nu\alpha\beta} F_{\mu\nu} Z_{\alpha\beta} - E \cdot \frac{\alpha_{\text{EM}} s_W^2}{8\pi c_W^2} \frac{a}{\langle\phi\rangle} \epsilon^{\mu\nu\alpha\beta} Z_{\mu\nu} Z_{\alpha\beta}, \end{aligned} \quad (4.18)$$

where s_W and c_W denote the sine and cosine of the Weinberg angle θ_W . Note that there are no anomalous W -couplings because we are only transforming right-handed fields, which are singlets under $SU(2)$. From here, it is customary to define

$$f_a \equiv \frac{1}{N_{\text{DW}}} \langle\phi\rangle, \quad (4.19)$$

where $N_{\text{DW}} \equiv 2N = 6$ is the domain-wall number counting the number of inequivalent vacua in the QCD-induced axion potential. The consequences of $N_{\text{DW}} \neq 1$ in our model are briefly discussed in section 7. The definition of f_a results in the conventional normalization of the axion couplings,

$$\begin{aligned} \mathcal{L} \supset & -\frac{\alpha_s}{16\pi f_a} a \epsilon^{\mu\nu\alpha\beta} G_{\mu\nu}^a G_{\alpha\beta}^a - \frac{E}{N} \cdot \frac{\alpha_{\text{EM}}}{16\pi f_a} a \epsilon^{\mu\nu\alpha\beta} F_{\mu\nu} F_{\alpha\beta} \\ & + \frac{E}{N} \cdot \frac{\alpha_{\text{EM}} s_W}{8\pi c_W} \frac{a}{f_a} \epsilon^{\mu\nu\alpha\beta} F_{\mu\nu} Z_{\alpha\beta} - \frac{E}{N} \cdot \frac{\alpha_{\text{EM}} s_W^2}{16\pi c_W^2} \frac{a}{f_a} \epsilon^{\mu\nu\alpha\beta} Z_{\mu\nu} Z_{\alpha\beta}. \end{aligned} \quad (4.20)$$

By comparison to the EFT Lagrangian defined in eqs. (2.1) and (2.3), we can identify c_{gg} , c_{WW} , $c_{\gamma\gamma}$, c_{ZZ} , and $c_{\gamma Z}$ as listed in table 1. If leptons are also charged under the PQ symmetry, an analogous axial rotation has to be performed in the leptonic sector, resulting in an additional contribution to the electromagnetic anomaly. In any case, eq. (4.20) is equally applicable with the corresponding value of E/N .

Axion couplings to light and heavy quarks. After absorbing the axionic phase of the mass matrix in the q_R fields, M becomes independent of a . Hence, U and S can also be chosen to be independent of a when proceeding with the subsequent steps in the diagonalization as described in eq. (4.14). After the diagonalization is completed, the

derivative axion-quark couplings induced by the transformation eq. (4.17) can be expressed in terms of the mass-diagonal fields as

$$\mathcal{L} \supset \frac{\langle \phi \rangle}{2} \begin{pmatrix} \bar{q}_R & \bar{F}_R \end{pmatrix} (\not{\partial} a) \Lambda^{-1} \begin{pmatrix} (\epsilon \epsilon')^2 \mathcal{A} \mathcal{B} \mathcal{A}^\dagger & \epsilon \epsilon' \mathcal{A} \mathcal{B} \\ \epsilon \epsilon' \mathcal{B} \mathcal{A}^\dagger & \mathcal{B} \end{pmatrix} \Lambda^{-1} \begin{pmatrix} q_R \\ F_R \end{pmatrix} \quad (4.21)$$

$$= \frac{1}{\langle \phi \rangle} \bar{q}_R (\not{\partial} a) q_R + \frac{\langle \phi \rangle}{2} \bar{F}_R (M_F^{-1} \mathcal{B} M_F^{-1}) (\not{\partial} a) F_R + \frac{v}{2} \epsilon' \bar{q}_R (M_q^{-1} \mathcal{A} \mathcal{B} M_F^{-1}) (\not{\partial} a) F_R + \text{h.c.}, \quad (4.22)$$

where \mathcal{A} and \mathcal{B} are coupling matrices which, to leading order in ϵ'' , are given by

$$\mathcal{A} = U_\delta^\dagger \alpha \lambda^{-1} U_\xi \quad \text{and} \quad \mathcal{B} = U_\xi^\dagger \beta \beta^\dagger U_\xi. \quad (4.23)$$

Note that the reason why the axion only couples to right-handed quarks is that we chose to only rotate these chiral components in eq. (4.17). This is clearly an arbitrary choice: we can perform an axion-dependent vector (non-axial) rotation in the quark field to include derivative couplings to left-handed quarks.

The field basis in which axions only couple derivatively is convenient for the calculation of loop processes as the one in section 5.1. But if we want to recover the effective terms in eq. (4.6), we have to rewrite the first term in eq. (4.22) by performing a rotation of the right-handed fields (neglecting interactions with the Higgs and CP-conserving anomalous axion terms),

$$\frac{1}{\langle \phi \rangle} \bar{q}_R (\not{\partial} a) q_R \rightarrow -\frac{a}{\langle \phi \rangle} \bar{q} M_q i \gamma_5 q. \quad (4.24)$$

Hence, we exactly arrive at purely pseudoscalar couplings to SM quarks which are proportional to the quark masses, as anticipated in eq. (4.6). Note that this implies that the right-handed flavour-diagonal coupling structure conserves CP, which can equally be shown by applying the CP transformation explicitly, see appendix A.

In contrast, applying the same steps to the second term in eq. (4.22) yields

$$\begin{aligned} & \frac{\langle \phi \rangle}{2} \bar{F}_R (M_F^{-1} \mathcal{B} M_F^{-1}) (\not{\partial} a) F_R \\ & \rightarrow -\frac{\langle \phi \rangle}{2} a \left[\bar{F} \frac{\mathcal{B} M_F^{-1} - M_F^{-1} \mathcal{B}}{2} i F + \bar{F} \frac{\mathcal{B} M_F^{-1} + M_F^{-1} \mathcal{B}}{2} i \gamma_5 F \right]. \end{aligned} \quad (4.25)$$

This demonstrates that the coupling of heavy F -quarks is not proportional to their masses, which is due to the fact that their mass is (mostly) generated by the spurion field Σ , whose VEV does not break the PQ symmetry. Furthermore, the interaction with F quarks is not of purely pseudoscalar form, which, as is discussed in section 5.2, has important implications for CP violation in our model.

Finally, for the last term in (4.22) we make use of the fact that q_R and F_R only differ in their respective masses but transform identically under the SM gauge group, and therefore

$$\partial_\mu (\bar{q}_{Ri} \gamma^\mu F_{Rj}) \rightarrow \bar{q}_i \frac{(M_q)_{ii} - (M_F)_{jj}}{2} i F_j + \bar{q}_i \frac{(M_q)_{ii} + (M_F)_{jj}}{2} i \gamma_5 F_j, \quad (4.26)$$

EFT coefficient	UV parameter	EFT model	EFT- ℓ model
f_a	$\langle\phi\rangle/(2N)$	$\langle\phi\rangle/6$	$\langle\phi\rangle/6$
c_{gg}	1	1	1
$c_{\gamma\gamma}$	E/N	5/3	8/3
c_{ZZ}	$\tan(\theta_W)^2 E/N$	$\tan(\theta_W)^2 5/3$	$\tan(\theta_W)^2 8/3$
$c_{\gamma Z}$	$-2 \tan(\theta_W) E/N$	$-2 \tan(\theta_W) 5/3$	$-2 \tan(\theta_W) 8/3$
c_{WW}	W/N	0	0
$(c_{q,R} - c_{q,L})$	$-(\chi_{q_R} - \chi_{q_L})/(2N)$	1/6 - 0	1/6 - 0
$(c_{\ell,R} - c_{\ell,L})$	$-(\chi_{\ell_R} - \chi_{\ell_L})/(2N)$	0	1/6 - 0

Table 1. Summary of couplings of the model introduced in the present section, expressed in the EFT basis introduced in eqs. (2.1) and (2.3). The general dependence on the high-energy parameters is shown, together with the explicit value for each of the two variations of the model, the one with and the one without tree-level couplings to leptons. In the three cases $c_{\gamma\gamma}$, $c_{\gamma Z}$, and c_{ZZ} , the coupling term should be written as in (4.20) with the electromagnetic fine structure constant α_{EM} . Further, the fermion coupling matrices in the last two rows are proportional to the unit matrix and $1/6 - 0$ indicates that only right-handed fields couple at tree-level.

where the indices i and j do not imply a sum but only a specific quark. This allows us to rewrite the last term in (4.22) as

$$\begin{aligned} & \frac{v}{2} \epsilon' \bar{q}_R (M_q^{-1} \mathcal{A} \mathcal{B} M_F^{-1}) (\not{\partial} a) F_R \\ \rightarrow & -\frac{v}{2} \epsilon' a \left[\bar{q} \frac{\mathcal{A} \mathcal{B} M_F^{-1} - M_q^{-1} \mathcal{A} \mathcal{B}}{2} i_F + \bar{q} \frac{\mathcal{A} \mathcal{B} M_F^{-1} + M_q^{-1} \mathcal{A} \mathcal{B}}{2} i_{\gamma_5} F \right]. \end{aligned} \quad (4.27)$$

Once more, the interaction contains both pseudoscalar and scalar parts which can in general be CP violating.

4.4 Low-energy effective couplings

To conclude this section, let us summarize how the model introduced above maps onto the generic axion EFT. In this context, it is important to note that the low-energy effective model does not contain any free parameter other than the axion decay constant and the mass of the heavy F -fermions, the latter only being relevant as a cut-off scale for loop processes (see the next section). All the couplings of the axion to SM particles are therefore fully determined once f_a is fixed. There is only one discrete choice, namely whether the SM leptons are charged under the PQ symmetry or not, leading to the two variations of the model. Table 1 summarizes the values that the EFT coefficients introduced in eqs. (2.1) and (2.3) take for both variations.

5 Flavour and CP effects in the EFT-inspired QCD axion model

In the previous section, we have introduced a QCD axion model with tree-level couplings to SM fermions and without any additional new states below the scale f_a other than the axion

itself. As a consequence, the model fully reproduces the phenomenological expectations based on the generic EFT description introduced in section 2, even when one-loop processes are involved. To confirm this, in this section, we delineate the main phenomenological features of this QCD axion, including a calculation of the $K^+ \rightarrow \pi^+ + a$ decay rate in the full model.

5.1 $K^+ \rightarrow \pi^+ + a$ decay rate

In the previous section, we have computed the effective interactions between the axion and the SM quarks, which are summarized in table 1. That said, our initial motivation for deriving the model in section 4 was to obtain an EFT-like model with no ambiguity in the log-enhanced FCNC contribution. If the interactions listed in table 1 were the only ones emerging in our UV model, there would be no additional contributions to the $K^+ \rightarrow \pi^+ + a$ amplitude at one loop and we would again be confronted with the same problem of UV divergences as before. Crucially, however, there are indeed further relevant interactions arising in our UV model that are not contained in table 1 and which render the loop computation finite. These operators are, as expected, related to the heavy F quarks and are hence not captured by the EFT approach.

To obtain these additional vertices, it is important to discuss further implications of the unitary transformation of the quark fields that we performed in eq. (4.14). As in the SM, this rotation generates new flavour-dependent interactions of neutral and charged hadronic currents with the weak gauge bosons. The detailed calculations are included in appendix C. The leading-order W interactions in the mass-diagonal basis can be written as

$$\mathcal{L} \supset \frac{-g}{\sqrt{2}} \begin{pmatrix} \bar{u}_L \\ \bar{d}_L \\ \bar{F}_L^u \\ \bar{F}_L^d \end{pmatrix}^T \gamma^\mu \left[W_\mu^+ \begin{pmatrix} 0 & V & 0 & -\epsilon\epsilon' V \mathcal{A}_d \\ 0 & 0 & 0 & 0 \\ 0 & -\epsilon\epsilon' \mathcal{A}_u^\dagger V & 0 & (\epsilon\epsilon')^2 \mathcal{A}_u^\dagger V \mathcal{A}_d \\ 0 & 0 & 0 & 0 \end{pmatrix} + W_\mu^- \begin{pmatrix} 0 & 0 & 0 & 0 \\ V^\dagger & 0 & -\epsilon\epsilon' V^\dagger \mathcal{A}_u & 0 \\ 0 & 0 & 0 & 0 \\ -\epsilon\epsilon' \mathcal{A}_d^\dagger V^\dagger & 0 & (\epsilon\epsilon')^2 \mathcal{A}_d^\dagger V^\dagger \mathcal{A}_u & 0 \end{pmatrix} \right] \begin{pmatrix} u_L \\ d_L \\ F_L^u \\ F_L^d \end{pmatrix}, \quad (5.1)$$

where we have reintroduced the labels for up- and down-type quarks as well as coupling matrices, and V denotes the CKM matrix as before. At leading order, V is given by $V = U_\delta^{u\dagger} U_\delta^d$ and turns out to be unitary up to terms of the order $(\epsilon\epsilon'')^2$. Similarly, the leading-order Z interactions become

$$\mathcal{L} \supset \begin{pmatrix} \bar{q} \\ \bar{F} \end{pmatrix}^T \gamma_\mu Z^\mu \frac{-g}{\cos(\theta_W)} \left[\pm \frac{1}{2} \begin{pmatrix} \mathbf{1} & -\epsilon\epsilon' \mathcal{A} \\ -\epsilon\epsilon' \mathcal{A}^\dagger & (\epsilon\epsilon')^2 \mathcal{A}^\dagger \mathcal{A} \end{pmatrix} P_L - Q \sin^2(\theta_W) \begin{pmatrix} \mathbf{1} & 0 \\ 0 & \mathbf{1} \end{pmatrix} \right] \begin{pmatrix} q \\ F \end{pmatrix}. \quad (5.2)$$

Here, the upper (lower) sign refers to up- (down-) type quarks, P_L denotes the projector onto left-handed fields and Q is the electromagnetic charge of each field. We find that only the left-handed coupling structure is affected by the transformation to the mass-diagonal basis.




Figure 3. Relevant one-loop contributions to the $s \rightarrow d + a$ process in our UV model. These consist of axion emissions through the quark couplings (left) as well as counterterm contributions (right). The symbol \times denotes a counterterm insertion.

With these additional interactions at hand, we can move forward to perform the computation of the loop processes shown in figure 3. The first diagram in figure 3 represents the sum of four individual contributions because each internal fermion propagator can either be a SM or an F quark. To calculate these diagrams, we make use of the axion-quark interactions as given in eq. (4.21) and the W couplings to hadronic currents in eq. (5.1). Expressing the interactions in this way means that the W boson and the axion only interact with left-handed and right-handed fields, respectively. Hence, we need mass insertions by each of the internal propagators. The corresponding factors of fermion masses are however cancelled by the inverse Λ matrices in the axion interaction. Putting everything together, it is easy to see that all four diagrams have the same flavour structure and are of the same order in ϵ and ϵ'' . We can write their combined contribution to the $s \rightarrow d + a$ amplitude in the compact form

$$\begin{aligned}
i\mathcal{M} = & i\frac{g^2}{2}(\epsilon\epsilon')^2\frac{\langle\phi\rangle}{2}\sum_{i,j,k,l}V_{di}^\dagger(\mathcal{A}_u)_{ij}(\mathcal{B}_u)_{jk}(\mathcal{A}_u^\dagger)_{kl}V_{ls}\int\frac{d^4k}{(2\pi)^4}\bar{u}_d(p_2)[\gamma_\nu\gamma_\rho(p_2-p_1)^\rho\gamma_\mu P_L]u_s(p_1) \\
& \times\left(\frac{1}{k^2-m_W^2}\left(g^{\mu\nu}-\frac{k^\mu k^\nu}{m_W^2}\right)\right) \\
& \times\left(\frac{1}{((p_2-k)^2+m_{Q_i}^2)((p_1-k)^2+m_{Q_i}^2)}-\frac{1}{((p_2-k)^2+m_{F_j}^2)((p_1-k)^2+m_{Q_i}^2)}\right. \\
& \left.-\frac{1}{((p_2-k)^2+m_{Q_i}^2)((p_1-k)^2+m_{F_k}^2)}+\frac{1}{((p_2-k)^2+m_{F_j}^2)((p_1-k)^2+m_{F_k}^2)}\right), \quad (5.3)
\end{aligned}$$

where m_Q refers to the up-type quarks. We evaluate the loop integral using **Package-X** [126] and under the simplifying assumption that all F quarks have equal masses. Only considering the leading-order terms in the down and strange fermion masses, we arrive at

$$i\mathcal{M} = h_{ds} \times \bar{u}_d(p_2) ((m_s - m_d) + (m_s + m_d)\gamma_5) u_s(p_1), \quad (5.4)$$

with

$$\begin{aligned}
h_{ds} = & -\frac{G_F}{16\sqrt{2}\pi^2}\frac{1}{\langle\phi\rangle} \\
& \times\sum_{q=u,c,t}V_{qd}^*V_{qs}m_q^2\left(\log\left(\frac{m_F^2}{m_q^2}\right)-\frac{2m_q^4-7m_q^2m_W^2+5m_W^4+3m_W^4\log\left(\frac{m_q^2}{m_W^2}\right)}{(m_q^2-m_W^2)^2}\right), \quad (5.5)
\end{aligned}$$

where we expanded the finite contributions in terms of $1/m_F$. This gives a contribution to $h_{ds}^S = h_{ds}$ as defined in the effective interaction Hamiltonian in (2.5). The logarithmically enhanced term is identical to the EFT result for $c_q = \frac{1}{6}$ and $\Lambda = m_F$. Note that the finite term is different from the EFT result as the loops with internal F -quarks also yield relevant contributions.

In principle, we also need to include diagrams exchanging the W boson in the first diagram in figure 3 with a Z boson. However, even if the couplings of the Z boson do have non-diagonal entries, they do not induce $s \rightarrow d + a$ at one-loop at the order in ϵ'' that we are considering. This is due to the matrix structure $\mathcal{A}\mathcal{B}\mathcal{A}^\dagger$ appearing in the relevant amplitudes, which is flavour-diagonal as it is identical to the matrix that determines the quark masses in eq. (4.15).

Finally, there are also counterterm contributions [127] from the renormalization of quark fields as depicted on the right of figure 3. Performing the calculations, however, one notices that the relevant contributions to eq. (5.4) cancel each other at linear order in the down- and strange-quark masses (see appendix D for more details). We can therefore also discard these counterterm diagrams and work with eq. (5.5) as our effective coefficient.

5.2 CP violation in axion interactions

From eq. (4.22), we observe that the axion has a right-handed coupling structure to all quarks in our theory. As discussed in appendix A, CP violation in flavour-violating couplings occurs when the corresponding coupling is complex valued. The axion interactions in eq. (4.22) can therefore induce CP violation since \mathcal{A} and/or \mathcal{B} can have imaginary entries. The fact that the interactions of the QCD axion can be CP violating may seem worrisome at first, since the QCD axion is introduced precisely to eliminate the CP violation in the strong sector. As shown in this section, however, the effects of this kind of CP violation on the electric dipole moment (EDM) of the neutron are comparable or even smaller to those induced by the phase of the CKM matrix, and are therefore not a threat to the axion solution of the strong CP problem. We point the reader to refs. [128, 129] for recent discussions on CP-violating axion interactions.

To see this, firstly note that this CP violation is of explicit nature, like the CKM phase in the SM, and not due to any spontaneous breaking. This can be understood considering the full UV theory in eq. (4.7). If the couplings α , β , and λ are real-valued, no CP-violating terms are generated before or after EW and PQ symmetry breaking. This is because in that case, U_δ and U_ξ appearing in eqs. (4.15) and (4.16) can be chosen to be orthogonal and real because the matrices that they diagonalize are real and symmetric to begin with. This then results in $\mathcal{A} = U_\delta^T \alpha \lambda^{-1} U_\xi$ and $\mathcal{B} = U_\xi^T \beta \beta^T U_\xi$ being real-valued as well. Hence, eq. (4.22) is CP conserving in this case. However, starting from complex-valued α , β , or λ couplings generically makes \mathcal{A} and/or \mathcal{B} complex valued as well. Therefore, the CP violation present in the axion interactions of eq. (4.22) is a reflection of potential explicit CP-breaking in the UV model. As a matter of fact, given that the CKM matrix, which is constructed from α , β , and λ , is experimentally confirmed to be complex valued, there is no reason to expect the coefficients \mathcal{A} and \mathcal{B} to be real without invoking an *ad hoc* cancellation.

The relevant question is therefore whether this CP violation is problematic, that is, whether it is in conflict with any observations. To answer this, we consider a very sensitive observable to CP violation: the neutron EDM d_n .

As shown in appendix E, one-loop processes do not generate quark EDMs in our UV model. Hence, a neutron EDM based on free quark EDMs may occur at the earliest at the two-loop level. Without delving into involved two-loop calculations, we instead derive a very conservative upper bound based on the expected scaling of the contributions. Assuming that at least two new physics vertices appear (such as internal emission and absorption of the axion), each of which carries a coupling suppression of f_a , we get the rough estimate

$$d_n^{\text{UV}} \lesssim \frac{e}{(16\pi^2)^2} \frac{m_n}{f_a^2} \approx 5 \cdot 10^{-32} e \cdot \text{cm}. \quad (5.6)$$

Here, m_n denotes the neutron mass, which we use as the characteristic scale, and we have inserted $f_a = 4 \cdot 10^6$ GeV as a lower-end value for the range of axion decay constants of interest. The above estimate is already very conservative in the sense that we have not included any suppression due to new physics contributions to W^\pm and/or Z gauge boson interactions and/or the heavy F -quark mass scale $\langle \Sigma \rangle$, since at least one of the three has to participate to generate a quark EDM. Moreover, we have not invoked any additional electric/weak coupling insertion or Yukawa/light quark mass suppression, the latter being expected to appear given the discussion above regarding the origin of the CP-violating interactions.

Interactions between the quark constituents within the neutron have also been shown to contribute to the neutron EDM (see ref. [130] for a review). That said, these effects are already below present sensitivity for purely SM-related interactions. This does not change in our UV model, as the new physics couplings/particles appearing would only lead to a stronger suppression. We can therefore safely neglect such contributions.

All in all and despite being very conservative, the estimate in eq. (5.6) is still much smaller than the current bounds [131] $d_n^{\text{exp}} \lesssim 10^{-26} e \text{cm}$. We therefore conclude that the resulting explicit CP violation is too small to affect near-future experiments measuring CP-sensitive observables and does not pose a threat to the axion solution to the strong CP problem in the model considered here.

6 Phenomenology of the model and discovery opportunities

Because our QCD axion model was constructed to reproduce the general features of the EFT setup, one may be sceptical that the model could possess any particular phenomenological feature that may serve as a handle to probe it. But interestingly, what at first seems like a lack of features does in fact have interesting consequences for experimental searches, which we now describe.

6.1 Astrophysical limits and helioscope searches

As any other QCD axion model, the one presented in this work is subject to constraints coming from its couplings to photons, electrons and nucleons. Figure 4 summarizes the bounds as a function of f_a and in the usual m_a - $g_{a\gamma\gamma}$ parameter space.

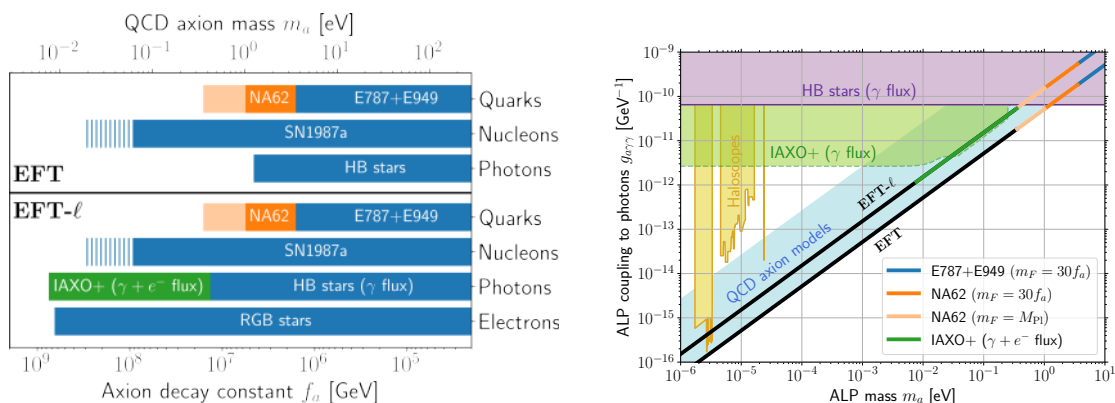


Figure 4. Constraints on the two variations of the model introduced in section 4, which we denote EFT- ℓ and EFT depending on whether tree-level axion-lepton couplings are present or not. On the left, we present the bounds as a function of the QCD axion decay constant. For the $K^+ \rightarrow \pi^+ + a$ constraints, the lighter shading of orange represents the range of possible values of the cutoff scale for the leading log, $\Lambda \in (30f_a, M_{Pl})$ as discussed in the main text. The vertical hatching indicates uncertainties in the supernova limit. On the right, we show the models predictions in the usual mass vs. photon coupling parameter space. The band of QCD axion models corresponds to $E/N \in (44/3, 5/3)$ as defined in [132]. Note that IAXO becomes particularly sensitive to the EFT- ℓ model due to solar axion production through axion-electron interactions. Hence, the predicted IAXO sensitivity in dark green is based on the interactions with both electrons and photons.

The axion-photon coupling is tested by a combination of laboratory experiments and astrophysical observations. The negative results from the solar axion search performed by CAST [10], together with the absence of exotic cooling in Horizontal Branch (HB) stars in globular clusters [133], constrain the photon coupling to be $g_{a\gamma\gamma} \lesssim 7 \times 10^{-11}$. This translates into a bound on f_a which is different for the leptonic and the non-leptonic models due to their different values of E/N (note that in the leptonic case, this bound only considers the axion flux generated by interactions with photons and not electrons; it is therefore overly conservative and only illustrative). Furthermore, the proposed successor of CAST, the IAXO solar telescope [11], is expected to improve this limit by over an order of magnitude in its upgraded version IAXO+. Unfortunately, the loss of sensitivity for axion masses above $\sim 10^{-2}$ eV, combined with an accidental cancellation¹¹ in the photon coupling, renders the non-leptonic model out of reach for IAXO+, as can be seen in the right panel of figure 4. The model with tree-level couplings to SM leptons, however, is within reach of the sensitivity forecast, as can be appreciated in the left panel of figure 4. The reason for this is the enhanced production of axions in the solar interior caused by the existence of a large electron coupling. Finally, haloscopes searches have an exquisite sensitivity to axions in the $1 - 100 \mu\text{eV}$ range, provided these particles were to make up the dark matter of the Universe. The currently excluded region, as compiled in [30, 69] and including new ADMX data [28, 29], is shaded in yellow in the right panel of figure 4.

¹¹The low-energy QCD axion-photon coupling is a combination of the model-dependent anomaly coefficient E/N and a model-independent piece coming from hadronic contributions [107].

Axions coupling to electrons induce additional cooling mechanisms in stars. Currently, the strongest limits come from observations of the brightness of the tip of the red-giant branch (RGB) in globular clusters, which exclude values of $f_a/c_e \geq 3.9 \times 10^9$ GeV, as derived in [134].¹² The axion-electron coupling is present at tree level in the leptonic variant of the model, and only loop induced (and therefore suppressed) in the non-leptonic one. As a consequence, the previously mentioned bound is only competitive for the former variation of the model, and suppressed by a loop factor [121, 139] $\sim \alpha_{\text{EM}}^2/(\pi^2) \log \sim 10^{-4}$ in the latter. Indeed, in the presence of tree-level leptonic couplings, RGB observations currently place the strongest constraints on the model. As mentioned above, these are only expected to be improved once IAXO has reached its full sensitivity.

Finally, we discuss the constraints arising from the effective coupling of axions to nucleons, which arises at low energies from the interactions with gluons and quarks. It is customary to define it in a way analogous to the other fermionic couplings as

$$\mathcal{L} \supset \frac{\partial_\mu a}{2f_a} \sum_{N=p,n} c_N \bar{N} \gamma^\mu \gamma_5 N. \tag{6.1}$$

The expression for the coefficients c_N in terms of the quark couplings can be found in [107], which shows good agreement with a recent reevaluation [140]. For the model at hand, they are $c_p = -0.39875$ and $c_n = 0.05125$. These couplings are best tested by studying their impact on the extreme dynamics of the proto-neutron star that forms in the course of a core-collapse supernova. Building on the seminal reference [12], the most recent limits are given in [83–85]. While there is overall agreement given the uncertainties, the bound of [84] on f_a is a factor of ~ 3.5 stronger than the most conservative ones in [83, 85]. In figure 4, we chose to use vertical hatching to showcase the spread in the different evaluations of the limit.

6.2 Constraints and opportunities from $K^+ \rightarrow \pi^+ + a$

Perhaps the most important feature of the model at hand is that quark-flavour violating transitions $q \rightarrow q' + a$, which are induced at the one-loop level, are enhanced by a large logarithm as discussed in the previous sections. This means that processes involving such transitions are particularly suitable for testing this QCD axion model. Among all such processes, the $s \rightarrow d + a$ decay is the one that offers the best experimental perspectives, thanks to the very precise measurements that can be performed at kaon facilities.

The predicted decay rate for $K^+ \rightarrow \pi^+ + a$ is given in eq. (5.4), where we can see that the leading-log piece matches the EFT result with the identification $\Lambda = m_F$. The corresponding branching ratio thus depends on the value of m_F . In principle, the only requirement for the mass of the heavy F -fermions is for them to be above the axion decay constant, $m_F \geq f_a$. However, given that SM Yukawa couplings are proportional to $\langle \phi \rangle / m_F$, it is possible to obtain some rough limits on the largest possible values of m_F that reproduce the SM quark masses. Assuming perturbativity of all dimensionless couplings, reproducing

¹²Slightly weaker limits come from measurements of the R parameter in globular clusters [135, 136] and from observations of white dwarfs, which exclude values of $f_a/c_e \geq 1.9 \times 10^9$ GeV, as derived in [137], although the data seems to prefer some amount of extra cooling compatible with a non-vanishing electron (and possibly photon) coupling [135, 138].

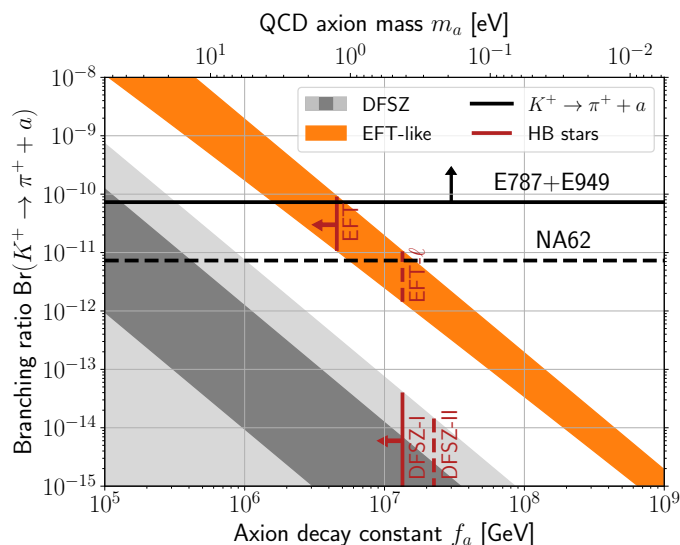


Figure 5. Comparison of the branching ratio of $K^+ \rightarrow \pi^+ + a$ predicted in the EFT-like model (orange band) and in the DFSZ model (grey bands). For the EFT-like model, we assume a universal quark coupling $c_q = 1/6$ and different values of the cutoff scale Λ (between $30f_a$ and M_{Pl} as discussed in the text). The DFSZ bands are obtained for a charged Higgs boson mass $m_{H^\pm} = 800$ GeV and $1 < \tan \beta < 5$ (darker band) or $0.25 < \tan \beta < 170$ (lighter band). The horizontal black lines show the existing constraint by the E787 and E949 experiments [116] (solid line) and the expected reach of NA62 [117] (dashed line). The vertical red lines describe the bounds set by HB stars, which depend on the photon coupling and therefore differ for the EFT/EFT- ℓ and DFSZ-I/II models.

the top quark Yukawa requires at least one of the F fermions to have a mass $m_{F_j} \lesssim 10^3 f_a$ for some j . However, the other F fermions could be much heavier, and without a complete construction of the 3-dimensional coupling matrices α , β , and γ , it is not possible to point down the contribution of each of the F fermions to $s \rightarrow d + a$ and therefore to the loop cutoff Λ . Therefore, and although it is reasonable to expect that $\Lambda \lesssim 10^3 f_a$, one cannot rule out the possibility of larger values.

Taking into account the discussion above, we can derive the bounds and prospects for this model in past and present kaon facilities. The $K^+ \rightarrow \pi^+ + a$ branching ratio as a function of f_a for different values of Λ corresponds to the orange lines in figure 5. A robust exclusion can be placed by assuming a conservative value for the cutoff, $\Lambda = 30 f_a$. Note that this choice is equivalent to $\epsilon'' \simeq 0.2$ and hence corresponds to a slight scale separation between $\langle \phi \rangle$ and $\langle \Sigma \rangle$. This choice may appear in tension with the assumption $\epsilon'' \ll 1$ employed in our discussion so far. However, as shown in appendix C for an explicit choice of coupling matrices, our expression for the relevant effective $s \rightarrow d + a$ coefficient derived in eq. (5.5) still holds to a very good degree in this case. With this, the E787 and E949 experiments rule out $f_a < 1.6 \times 10^6$ GeV, and NA62 could push this up to $f_a < 5.6 \times 10^6$ GeV (see the left panel of figure 4, dark orange) assuming an order of magnitude better sensitivity in the branching ratio. That said, given that Λ could be much larger, the discovery potential for these experiments extends somewhat further. Only in

this way, by identifying m_F as the scale appearing in the logarithm and treating it as an independent parameter from the PQ breaking scale $\langle\phi\rangle$, can the NA62 sensitivity go beyond the robust limits set by the cooling of HB stars for the EFT scenario. For the most extreme value of $\Lambda = M_{\text{pl}}$, the potential reach of NA62 extends to $f_a \lesssim 1.6 \times 10^7 \text{ GeV}$, which as can be seen in the left panel of figure 4 (light orange) comes closer to the SN1987a limit.

7 Conclusions and discussion

Effective Field Theory tools are a cornerstone of studies of the QCD axion and general axion-like particles. However, the validity of an EFT is limited by the existence of a cutoff. This is evident in log-divergent loop processes, where the cutoff explicitly appears in the final result. It is tempting to infer the cutoff scale from the scale set by the dimensionful coupling constants of the higher dimensional operators in the axion EFT. Yet, one may wonder if this choice is validated by embeddings into simple axion models with renormalizable couplings. In this work, we have compared calculations made using an EFT setup with those obtained using complete QCD axion models. Our results show that matching EFT predictions to full models is far from trivial when loop processes are involved.

Owing to its experimental relevance, we have chosen to focus on the rare decay $K^+ \rightarrow \pi^+ + a$ for our study. We have shown that the $K^+ \rightarrow \pi^+ + a$ decay rate in the most popular QCD axion models can differ by orders of magnitude from naive EFT predictions, similar to what was found in [87] based on effective models that respect electroweak symmetry. The origin of this discrepancy lies in the UV sensitivity that is induced by the logarithmic divergence of this loop-induced decay. For DFSZ models, it is the existence of new degrees of freedom much below the PQ symmetry-breaking scale which violates the assumptions on which the EFT loop calculation is based. On the other hand, for KSVZ (or hadronic) QCD axion models, which do not have extra degrees of freedom below the PQ scale, suitable tree-level couplings to SM fermions are absent.

The above considerations appear to question the existence of a complete axion model which can successfully reproduce the EFT result, i.e. the large logarithmic enhancement found in the EFT calculation for loop-induced rare decays involving an axion. We have addressed this issue by explicitly constructing such a model, whose low-energy dynamics are dictated by the Lagrangian in eq. (4.1). Crucially, this new model allows for tree-level axion-SM fermion couplings without introducing any extra degree of freedom below the PQ symmetry-breaking scale other than the axion itself.

While very minimal at energies below f_a , the EFT-inspired QCD axion model constructed in this work presents rich dynamics above the PQ scale. In addition to the PQ complex scalar, three generations of heavy coloured fermions have to be introduced in order to obtain a renormalizable model. Through mixing with SM quarks, these heavy fermions induce modifications of the low-energy axion couplings. Ultimately, it is the presence of these additional states which tames the logarithmic divergence of the $K^+ \rightarrow \pi^+ + a$ rate, therefore providing a physically-motivated cutoff as opposed to the artificial one that is employed in the EFT.

The implementation that we have followed in this work is only one of the possibilities to UV complete the low-energy Lagrangian eq. (4.1). Some other scenarios are delineated in appendix B. These may feature qualitatively different dynamics for the interactions of the QCD axion with the SM fields. One example that is briefly described in appendix C is the alternative that the mass of the heavy quarks lies close (or even below) the PQ-breaking scale. In this case, the QCD axion generically enjoys tree-level flavour-violating interactions with the SM quarks, which significantly change the flavour phenomenology (see, e.g. [71]). Other interesting possibilities, like potential CP-violating axion interactions or connections between the axion dynamics and the flavour puzzle are interesting avenues of further research. Finally, the large logarithms appearing together with the top-quark Yukawa suggests upgrading our loop calculations with RG improvements along the lines of [87–89] to achieve better accuracy in the phenomenological limits.

The EFT-inspired QCD axion model has a rich phenomenology, the main features of which are summarized in figure 4. As is also the case in other axion models, astrophysical observables currently place the strongest constraints on axions with sub-eV masses. In particular, the duration of the neutrino burst in the supernova 1987a excludes decay constants $f_a \lesssim 10^8$ GeV. However, owing to the logarithmic enhancement of the $K^+ \rightarrow \pi^+ + a$ rate, $K^+ \rightarrow \pi^+ + \text{inv}$ searches at the NA62 experiment have the potential to supersede all astrophysical bounds except for the SN1987a one. Given the uncertainties in the modelling of axion emission in supernovae,¹³ a confirmation of the exclusion from a terrestrial experiment is highly desirable. Additionally, the presence of additional colour-charged fermions with PQ charge could act to decrease the value of the QCD anomaly coefficient N , in which case the SN1987a limit would be weakened compared to the NA62 one.

The possibility to couple the axion to the SM leptons in a similar way as we have done for quarks offers even more exciting possibilities. Due to the particular balance between the photon and the electron couplings in our model, the IAXO helioscope (in its upgraded version IAXO+) is expected to become the most sensitive probe of sub-eV mass QCD axions, with a reach close to 10^9 GeV in f_a (corresponding to a mass as small as 0.01 eV). This would supersede all astrophysical limits, including the previously mentioned SN1987a one and the one derived from observations of the brightness of the tip of the red giant branch in globular clusters.

The dark matter phenomenology of the investigated models is similar to those in the DFSZ and KSVZ benchmark models, with two notable differences. One is that the photon coupling can be quite small due to an electromagnetic anomaly coefficient of 5/3. Second, in the models that we have investigated the domain-wall number is greater than one, favouring the scenario where the PQ symmetry is broken during inflation and not restored afterwards. Alternatively, one may have an additional explicit breaking of the discrete Z_N symmetry. While this can be useful to allow for the correct dark matter abundance at smaller values of the PQ scale f_a [144], potentially accessible with experiments such as NA62, it may also lead to a certain degree of tuning.

¹³Some authors [141] have proposed an alternative scenario for the SN1987a explosion, which would render the SN bound on axions invalid. However, [142, 143] argue that recent hints towards the existence of a neutron star remnant in the site of the explosion favour the standard hypothesis.

The findings of this work showcase the necessity of developing a detailed understanding of the strengths and limitations of EFT tools in axion studies. We have seen that it is far from trivial to construct a UV embedding of an axion EFT coupled to fermions that agrees with the leading log predictions for relevant experimental processes. More generally, this leads to the question of whether a generic axion EFT with a cutoff given by the (inverse of) the dimension-5 coupling constants can be obtained from a UV model. While our explicit construction answers this question in an affirmative manner for the specific case at hand, it also highlights potential issues such as the quite complicated structure of the model, the potential appearance of additional flavour-changing couplings as well as potentially non-perturbative Yukawas being required to reproduce the top-quark mass.¹⁴ It therefore remains an intriguing question whether even mild extra assumptions may severely restrict the space of axions EFTs that are the low-energy manifestations of reasonable UV models, possibly giving rise to an *ALPine swampland*.

Acknowledgments

We thank Susanne Westhoff for helpful discussions. GA acknowledges support by a “la Caixa” postgraduate fellowship from the Fundación “la Caixa”, by a McGill Trottier Chair Astrophysics Postdoctoral Fellowship, and by NSERC (Natural Sciences and Engineering Research Council, Canada). FE and FK are funded by the Deutsche Forschungsgemeinschaft (DFG) through the Emmy Noether Grant No. KA 4662/1-1. LT is funded by the Graduiertenkolleg *Particle physics beyond the Standard Model* (GRK 1940).

A CP violation in pseudoscalar interactions

In this appendix, we investigate under what conditions CP violation arises in the couplings studied in this work. In particular, we focus on axion-quark, W -quark, and Z -quark interactions, where quark generally refers to a quark of any type, i.e. SM or BSM. The flavour-diagonal structure of photon and gluon interactions is not altered in our UV model compared to the SM and hence does not require any further comments.

Axion-quark interactions. We begin with derivative axion-quark interactions, differentiating between flavour-violating and flavour-conserving operators. The latter can in full generality be written as

$$(\partial_\mu a) \bar{q} \gamma^\mu (h^S + h^P \gamma_5) q. \tag{A.1}$$

Hermiticity of this operator restricts h^S and h^P to be real-valued. A CP transformation leaves (A.1) unchanged and hence this interaction does never violate CP. On the other hand, the flavour-violating version of this operator and its hermitian conjugate read

$$(\partial_\mu a) \bar{q} \gamma^\mu (h^S + h^P \gamma_5) q' + (\partial_\mu a) \bar{q}' \gamma^\mu \left((h^S)^* + (h^P)^* \gamma_5 \right) q. \tag{A.2}$$

¹⁴Problems in obtaining the top-quark mass while retaining perturbativity were already noted in [87].

Applying the CP operation on the first term, we obtain

$$(\partial_\mu a) \bar{q} \gamma^\mu (h^S + h^P \gamma_5) q' \xrightarrow{\text{CP}} (\partial_\mu a) \bar{q}' \gamma^\mu (h^S + h^P \gamma_5) q, \quad (\text{A.3})$$

and hence CP violation is closely related to the imaginary parts of h^S and h^P . Importantly, CP violation does not arise even if the axion simultaneously enjoys scalar and pseudoscalar couplings in the derivative basis, as long as both coupling constants remain real. For example, in the EFT scenario in eqs. (2.5) and (2.8), the CP violation is fully determined by the imaginary part of the CKM matrix elements.

Gauge boson-quark interactions. We start with the Z bosons-quark couplings. For flavour-conserving interactions, we have

$$\bar{q} \gamma_\mu (V + A \gamma_5) Z^\mu q, \quad (\text{A.4})$$

where hermiticity requires both couplings V and A to be real-valued. As a consequence, invariance under a CP transformation is automatically guaranteed. In contrast, flavour-violating couplings involve

$$\bar{q} \gamma_\mu (V + A \gamma_5) Z^\mu q' + \bar{q}' \gamma_\mu (V^* + A^* \gamma_5) Z^\mu q. \quad (\text{A.5})$$

A CP transformation on the first operator gives

$$\bar{q} \gamma_\mu (V + A \gamma_5) Z^\mu q' \xrightarrow{\text{CP}} \bar{q}' \gamma_\mu (V + A \gamma_5) Z^\mu q, \quad (\text{A.6})$$

and hence CP violation is again closely related to complex-valued couplings, i.e. phases of V and/or A . In our case the coupling structure for flavour-violating interactions of the Z bosons is always left-handed (see eq. (C.47)) and hence we have $V = -A$. CP violation therefore also amounts here to a global phase attached to the operator $\bar{q} \gamma_\mu Z^\mu P_L q'$.

The situation of W bosons is simpler given that in general only left-handed fields couple to it. The coupling structure are always of the type

$$\bar{q}(V - A) \gamma_\mu W^\mu P_L q' + \text{h.c.} \quad (\text{A.7})$$

As is familiar from the CKM matrix in the SM, CP violation manifests itself when $(V - A)$ is complex-valued.

B Variations of the model

Instead of the specific UV completion in eq. (4.7), we could have allowed either or both of the up- and down-type fields to couple to Φ^* instead of Φ . In total, there are four possible ways to choose between Φ and Φ^* . The most general Lagrangian reads

$$\begin{aligned} \mathcal{L} \supset & -\alpha^u \bar{Q}_L \tilde{H} F_R^u - \beta^u \bar{F}_L^u \Phi^{(*)} u_R + \text{h.c.} \\ & -\alpha^d \bar{Q}_L H F_R^d - \beta^d \bar{F}_L^d \Phi^{(*)} d_R + \text{h.c.} \end{aligned} \quad (\text{B.1})$$

Since Φ is a SM scalar, this choice does not influence the charges of any fields under the SM gauge groups. However, each case results in different relations between PQ charges of the different fermions. The terms in eq. (B.1) enforce the conditions

$$\chi_{Q_L} - \chi_{F_R^u} = \chi_{Q_L} - \chi_{F_R^d} = 0, \tag{B.2}$$

$$\chi_{F_L^u} - \chi_{u_R} = \pm 1, \tag{B.3}$$

$$\chi_{F_L^d} - \chi_{d_R} = \pm 1, \tag{B.4}$$

where the upper sign corresponds to Φ and the lower sign to Φ^* in the Lagrangian (B.1). Because all q and F quarks are in the fundamental representation of SU(3), the sum of the four equations above appears in the anomaly coefficient,

$$N = \sum_f (\chi_{f_L} - \chi_{f_R}) T(R_f). \tag{B.5}$$

Hence, $|N| = 3$ as long as we choose either Φ or Φ^* for both up- and down-type quarks. Otherwise, the two contributions cancel, there is no QCD anomaly and a is not a QCD axion. We therefore do not further pursue that last possibility. Therefore, PQ charges of up- and down-type fields are necessarily identical, and we drop the corresponding labels in what follows.

After this, the only other choice is the PQ charge assignment of the F quarks. This determines the possible origin of the F -quark masses, which can either originate from the PQ-scalar Φ or from the VEV of an additional spurion field Σ without PQ charge. The combination of the two possible charge assignments of q and F quarks results in four different models:

1. Use Φ in (B.1) and generate F masses from $\langle \Sigma \rangle$. The additional mass term is

$$\mathcal{L} \supset -\lambda \frac{\langle \Sigma \rangle}{\sqrt{2}} \bar{F}_L F_R + \text{h.c.}, \tag{B.6}$$

which implies

$$\chi_{F_L} - \chi_{F_R} = 0 \quad \Rightarrow \quad \chi_{Q_L} - \chi_{q_R} = 1. \tag{B.7}$$

F is vector-like with respect to the PQ symmetry and does not contribute to the anomaly. This is exactly the model which is analyzed in great detail in sections 4 and 5.

2. Use Φ^* in (B.1) and generate F masses from $\langle \Sigma \rangle$. Using the same mass term as in (B.6), we get

$$\chi_{F_L} - \chi_{F_R} = 0 \quad \Rightarrow \quad \chi_{Q_L} - \chi_{q_R} = -1. \tag{B.8}$$

This option is of course equivalent to the first one after a redefinition of $\Phi \leftrightarrow \Phi^*$.

3. Use Φ in (B.1) and generate F masses from $\langle \Phi \rangle$. The mass term of F quarks in this case is given by

$$\mathcal{L} \supset -\lambda \frac{\langle \phi \rangle}{\sqrt{2}} \bar{F}_L F_R + \text{h.c.}. \tag{B.9}$$

And consequently,

$$\chi_{F_L} - \chi_{F_R} = 1 \quad \Rightarrow \quad \chi_{Q_L} - \chi_{q_R} = 0. \quad (\text{B.10})$$

Only F carries an axial PQ charge, which means that no symmetry forbids a Yukawa term as in the SM,

$$\mathcal{L} \supset -\gamma^u \bar{Q}_L \tilde{H} u_R - \gamma^d \bar{Q}_L H d_R + \text{h.c.} \quad (\text{B.11})$$

The existence of these couplings slightly modifies the diagonalization of the mass matrix. Even though the q fields are not charged under PQ in the UV model, they inherit axion couplings through the mixing with the heavy quarks during the diagonalization. These are however parametrically suppressed by $\epsilon^2 = \frac{v^2}{\langle \phi \rangle^2}$.

4. Use Φ^* in (B.1) and generate F masses from $\langle \Phi \rangle$. Using the same mass term as in (B.9), we get

$$\chi_{F_L} - \chi_{F_R} = 1 \quad \Rightarrow \quad \chi_{Q_L} - \chi_{q_R} = -2. \quad (\text{B.12})$$

Both q and F quarks have axial PQ charges and therefore tree-level axion couplings in this variation of the model.

In each of the cases above, all axial charges of all combinations of SM and F quarks are fixed. The only remaining freedom is a shift of all charges by an arbitrary constant, which can be used to fix one vector-like charge.

Note that it only makes sense to work with the effective Lagrangian eq. (4.1) in the first (or the equivalent second model) when $\langle \phi \rangle \ll m_{F_i}$. Otherwise, the F fields cannot be integrated out at any scale above PQ symmetry breaking as is done in the main text. In models 3 and 4, a full diagonalization including all q and F quarks along the lines of what is done in appendix C is always necessary.

C Mass diagonalization and axion interactions

This appendix contains the full diagonalization procedure of the quark mass matrix as well as the resulting couplings of the axion and light and heavy quarks. For convenience, this includes some repetitions of intermediate steps and results which are also included in section 4.3.

Absorbing the axion field into the quarks. To have a field basis in which all fields are mass eigenstates, we have to diagonalize the mass matrix given by

$$M = \frac{\langle \Sigma \rangle}{\sqrt{2}} \begin{pmatrix} 0 & \epsilon \epsilon' \alpha \\ \epsilon' \beta e^{ia/\langle \Phi \rangle} & \lambda \end{pmatrix}, \quad (\text{C.1})$$

with the two expansion parameters

$$\epsilon = \frac{v}{\langle \phi \rangle} \quad \text{and} \quad \epsilon' = \frac{\langle \phi \rangle}{\langle \Sigma \rangle}. \quad (\text{C.2})$$

As in the main text, we split the diagonalization into two parts, of which only the first one depends on the axion field. We start by absorbing the axion dependence of the quark mass matrix into the right-handed quarks,

$$u_R \rightarrow e^{-\frac{ia}{\langle\phi\rangle}} u_R, \quad d_R \rightarrow e^{-\frac{ia}{\langle\phi\rangle}} d_R. \quad (\text{C.3})$$

This transformation removes the axion field a from M in eq. (C.1) while the quark kinetic terms generate derivative couplings of the axion to right-handed quarks,

$$\bar{q}_R i \not{\partial} q_R \rightarrow \bar{q}_R i \not{\partial} q_R + \frac{\partial_\mu a}{\langle\phi\rangle} \bar{q}_R \gamma^\mu q_R, \quad (\text{C.4})$$

where q stands for both up- and down-type fields. The path integral measure is not invariant under this transformation and anomalous interaction terms between the axion and gauge bosons arise. These are [90, 91, 123–125]

$$\begin{aligned} \mathcal{L} \supset & -2N \cdot \frac{\alpha_s}{16\pi \langle\phi\rangle} a \epsilon^{\mu\nu\alpha\beta} G_{\mu\nu}^a G_{\alpha\beta}^a - E \cdot \frac{\alpha_{\text{EM}}}{8\pi \langle\phi\rangle} a \epsilon^{\mu\nu\alpha\beta} F_{\mu\nu} F_{\alpha\beta} \\ & + E \cdot \frac{\alpha_{\text{EM}} s_W}{4\pi c_W \langle\phi\rangle} a \epsilon^{\mu\nu\alpha\beta} F_{\mu\nu} Z_{\alpha\beta} - E \cdot \frac{\alpha_{\text{EM}} s_W^2}{8\pi c_W^2 \langle\phi\rangle} a \epsilon^{\mu\nu\alpha\beta} Z_{\mu\nu} Z_{\alpha\beta}. \end{aligned} \quad (\text{C.5})$$

With eq. (C.4) we have chosen to apply axion-dependent rotations only to right-handed fields which do not couple to SU(2) gauge fields. Therefore, W -couplings are absent. N and E are the anomaly coefficients defined as

$$N = \sum_f (\chi_{f_L} - \chi_{f_R}) T(R_f) = 3, \quad (\text{C.6})$$

$$E = \sum_f (\chi_{f_L} - \chi_{f_R}) Q_f^2. \quad (\text{C.7})$$

$T(R_f)$ is the Dynkin index of the SU(3) representation and Q_f are the electric charges. In our model, $N = 3$ and $E = 5$ or 8 depending on whether the leptons are also charged under PQ. We normalize the axion gluon coupling by defining

$$f_a \equiv \frac{1}{N_{\text{DW}}} \langle\phi\rangle, \quad (\text{C.8})$$

where $N_{\text{DW}} \equiv 2N = 6$ is the domain wall number counting the number of inequivalent vacua in the QCD-induced axion potential. This means that we can write the gauge-boson coupling terms as

$$\begin{aligned} \mathcal{L} \supset & -\frac{\alpha_s}{16\pi f_a} a \epsilon^{\mu\nu\alpha\beta} G_{\mu\nu}^a G_{\alpha\beta}^a - \frac{E}{N} \cdot \frac{\alpha_{\text{EM}}}{16\pi f_a} a \epsilon^{\mu\nu\alpha\beta} F_{\mu\nu} F_{\alpha\beta} \\ & + \frac{E}{N} \cdot \frac{\alpha_{\text{EM}} s_W}{8\pi c_W f_a} a \epsilon^{\mu\nu\alpha\beta} F_{\mu\nu} Z_{\alpha\beta} - \frac{E}{N} \cdot \frac{\alpha_{\text{EM}} s_W^2}{16\pi c_W^2 f_a} a \epsilon^{\mu\nu\alpha\beta} Z_{\mu\nu} Z_{\alpha\beta}. \end{aligned} \quad (\text{C.9})$$

Comparing to the EFT setup defined in eq. (2.1), we can identify $c_{gg} = 1$, $c_{WW} = 0$, $c_{\gamma\gamma} = E/N$, $c_{ZZ} = \tan(\theta_W)^2 E/N$ and $c_{\gamma Z} = -2 \tan(\theta_W) E/N$.

We have eliminated the a dependence of the mass matrix by the transformation in eq. (C.3). The subsequent steps in the diagonalization procedure also contain axial transformations of the quark fields, which are however not axion dependent and only result in the usual shift of the theta angle of QCD θ_{QCD}

$$\theta_{\text{QCD}} \rightarrow \theta_{\text{QCD}} + \arg(\det(M|_{a=0})), \quad (\text{C.10})$$

therefore simply displacing the location of the minimum of the axion potential.

Mass diagonalization. We continue with fully diagonalizing the matrices M^u and M^d by unitary transformations of left- and right handed fields. As before, we drop the labels u/d . At every point, we consider all three generations of SM and F quarks even though the indices are also omitted.

Because MM^\dagger is hermitian, it can be diagonalized by a unitary matrix U ,

$$U^\dagger MM^\dagger U = \Lambda^2, \quad (\text{C.11})$$

where Λ^2 is a diagonal matrix with only real positive eigenvalues. We then define the unitary matrix $S = M^\dagger U \Lambda^{-1}$ and a unitary transformation of fields

$$\begin{pmatrix} q_L \\ F_L \end{pmatrix} \rightarrow U \begin{pmatrix} q_L \\ F_L \end{pmatrix}, \quad (\text{C.12})$$

$$\begin{pmatrix} q_R \\ F_R \end{pmatrix} \rightarrow S \begin{pmatrix} q_R \\ F_R \end{pmatrix}. \quad (\text{C.13})$$

Here, q can be either u or d . This transformation diagonalizes the mass matrix

$$U^\dagger M S = U^\dagger M M^\dagger U \Lambda^{-1} = \Lambda. \quad (\text{C.14})$$

We now need to perturbatively find U , which diagonalizes

$$M M^\dagger = \frac{\langle \Sigma \rangle^2}{2} \begin{pmatrix} \epsilon^2 \epsilon'^2 (\alpha \alpha^\dagger) & \epsilon \epsilon' (\alpha \lambda^\dagger) \\ \epsilon \epsilon' (\lambda \alpha^\dagger) & (\lambda \lambda^\dagger + \epsilon'^2 \beta \beta^\dagger) \end{pmatrix} \equiv \frac{\langle \Sigma \rangle^2}{2} \begin{pmatrix} (\epsilon \epsilon')^2 \delta & \epsilon \epsilon' \mu \\ \epsilon \epsilon' \mu^\dagger & \xi \end{pmatrix}. \quad (\text{C.15})$$

In the last step we have defined the matrices δ , μ and ξ for notational convenience, of which δ and ξ are hermitian. To quadratic order in ϵ , U is given by

$$U = \begin{pmatrix} (-1 + \frac{(\epsilon \epsilon')^2}{2} \mu \xi^{-2} \mu^\dagger) U_\delta & \epsilon \epsilon' \mu \xi^{-1} U_\xi \\ \epsilon \epsilon' \xi^{-1} \mu^\dagger U_\delta & (1 - \frac{(\epsilon \epsilon')^2}{2} \xi^{-1} \mu^\dagger \mu \xi^{-1}) U_\xi \end{pmatrix} + \mathcal{O}(\epsilon^3), \quad (\text{C.16})$$

where the unitary matrices U_δ and U_ξ are defined by the property that they diagonalize hermitian matrices to give the q and F masses. Respectively,

$$M_q^2 = \text{diag}(m_{q_1}^2, m_{q_2}^2, m_{q_3}^2) = \left[U_\delta^\dagger (\delta - \mu \xi^{-1} \mu^\dagger) U_\delta (\epsilon \epsilon')^2 + \mathcal{O}(\epsilon^3) \right] \frac{\langle \Sigma \rangle^2}{2}, \quad (\text{C.17})$$

$$M_F^2 = \text{diag}(m_{F_1}^2, m_{F_2}^2, m_{F_3}^2) = \left[U_\xi^\dagger \left(\xi + \frac{(\epsilon \epsilon')^2}{2} (\mu^\dagger \mu \xi^{-1} + \xi^{-1} \mu^\dagger \mu) \right) U_\xi + \mathcal{O}(\epsilon^3) \right] \frac{\langle \Sigma \rangle^2}{2}. \quad (\text{C.18})$$

In the region of interest for NA62, this expansion in ϵ can be done safely, as $\epsilon \lesssim 10^{-4}$.

In order to map onto the effective description and to avoid sizable corrections to the effective description in (4.1), we need to ensure that $m_{F_i} \gg \langle \phi \rangle$ such that the F quarks can be integrated out at some scale above the PQ one. In other words, we need the scale separation between the mass of the F quarks and the PQ scale to be sufficiently large. This condition can be written as

$$\min_i m_{F_i}^2 = \min \text{eig}(\xi) \langle \Sigma \rangle^2 \gg \langle \phi \rangle^2 \quad (\text{C.19})$$

$$\Rightarrow 1 \gtrsim \min \text{eig}(\xi) = \min \text{eig}(\lambda\lambda^\dagger + \epsilon'^2 \beta\beta^\dagger) \gg \epsilon'^2 \quad (\text{C.20})$$

$$\Rightarrow 1 \gtrsim \min \text{eig}(\lambda\lambda^\dagger) \gg \epsilon'^2. \quad (\text{C.21})$$

Here, $\min \text{eig}$ denotes the smallest eigenvalue of a matrix. In the second and third lines, the size of the eigenvalues are constrained by perturbativity. We see that $\epsilon' \ll 1$ is only a necessary condition for the F quarks to be much heavier than the PQ scale, while the last line is a sufficient condition. We define one more expansion parameter ϵ'' as

$$\epsilon''^2 = \frac{\epsilon'^2}{\min \text{eig}(\lambda\lambda^\dagger)} \simeq \frac{\langle \phi \rangle^2}{\min_i m_{F_i}^2}. \quad (\text{C.22})$$

When $\epsilon'' \ll 1$, we can expand ξ^{-1} as

$$\xi^{-1} = (\lambda\lambda^\dagger + \epsilon'^2 \beta\beta^\dagger)^{-1} \quad (\text{C.23})$$

$$= \lambda^{\dagger-1} (1 + \epsilon'^2 \lambda^{-1} \beta\beta^\dagger \lambda^{\dagger-1})^{-1} \lambda^{-1} \quad (\text{C.24})$$

$$= \lambda^{\dagger-1} \sum_n (-\epsilon'^2 \lambda^{-1} \beta\beta^\dagger \lambda^{\dagger-1})^n \lambda^{-1} \quad (\text{C.25})$$

$$= \lambda^{\dagger-1} (1 - \epsilon'^2 \lambda^{-1} \beta\beta^\dagger \lambda^{\dagger-1} + \mathcal{O}(\epsilon''^4)) \lambda^{-1}, \quad (\text{C.26})$$

where we have used a Neumann series from the second to the third line and assumed that eigenvalues of $\beta\beta^\dagger$ are at most of order 1. Inserting the leading-order result in ϵ'' into (C.17), we find for the light quark masses

$$M_q^2 = \text{diag}(m_{q_1}^2, m_{q_2}^2, m_{q_3}^2) \simeq U_\delta^\dagger \alpha \lambda^{-1} \beta\beta^\dagger \lambda^{\dagger-1} \alpha^\dagger U_\delta \epsilon^2 \epsilon'^4 \frac{\langle \Sigma \rangle^2}{2} \simeq \mathcal{A} \mathcal{B} \mathcal{A}^\dagger \epsilon^2 \epsilon'^4 \frac{\langle \Sigma \rangle^2}{2}. \quad (\text{C.27})$$

In the last step, we have defined the coupling matrices

$$\mathcal{A} = U_\delta^\dagger \mu \xi^{-1} U_\xi = U_\delta^\dagger \alpha \lambda^{-1} U_\xi + \mathcal{O}(\epsilon''^2) \quad \text{and} \quad \mathcal{B} = U_\xi^\dagger \beta\beta^\dagger U_\xi. \quad (\text{C.28})$$

Axion-quark interactions. Because the axion field was absorbed entirely into the right-handed fields, the derivative terms with left-handed quark fields are not affected by the unitary transformation U . However, the derivative axion coupling to right-handed fields as in eq. (C.4) does not transform trivially,

$$\frac{\partial_\mu a}{\langle \phi \rangle} (\bar{q}_R \bar{F}_R) \gamma^\mu \begin{pmatrix} \mathbb{1} & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} q_R \\ F_R \end{pmatrix} \rightarrow \frac{\partial_\mu a}{\langle \phi \rangle} (\bar{q}_R \bar{F}_R) \gamma^\mu S^\dagger \begin{pmatrix} \mathbb{1} & 0 \\ 0 & 0 \end{pmatrix} S \begin{pmatrix} q_R \\ F_R \end{pmatrix}. \quad (\text{C.29})$$

To leading order in ϵ , we obtain for S

$$S = M^\dagger U \Lambda^{-1} \tag{C.30}$$

$$= \frac{\langle \Sigma \rangle}{\sqrt{2}} \left[\begin{pmatrix} \epsilon \epsilon'^2 \beta^\dagger \xi^{-1} \mu^\dagger U_\delta & \epsilon' \beta^\dagger (1 - \frac{(\epsilon \epsilon')^2}{2} \xi^{-1} \mu^\dagger \mu \xi^{-1}) U_\xi \\ \epsilon \epsilon' (\lambda^\dagger \xi^{-1} \mu^\dagger - \alpha^\dagger) U_\delta & (\lambda^\dagger + (\epsilon \epsilon')^2 (\alpha^\dagger \mu \xi^{-1} - \frac{1}{2} \lambda^\dagger \xi^{-1} \mu^\dagger \mu \xi^{-1})) U_\xi \end{pmatrix} + \mathcal{O}(\epsilon^3) \right] \Lambda^{-1}, \tag{C.31}$$

from which we can determine the relevant coupling matrix to leading order in ϵ as

$$\frac{\not{\partial} a}{\langle \phi \rangle} S^\dagger \begin{pmatrix} \mathbf{1} & 0 \\ 0 & 0 \end{pmatrix} S = \frac{\not{\partial} a}{\langle \phi \rangle} \Lambda^{-1} \frac{\langle \Sigma \rangle^2}{2} \begin{pmatrix} \epsilon^2 \epsilon'^4 U_\delta^\dagger \mu \xi^{-1} \beta \beta^\dagger \xi^{-1} \mu^\dagger U_\delta & \epsilon \epsilon'^3 U_\delta^\dagger \mu \xi^{-1} \beta \beta^\dagger U_\xi \\ \epsilon \epsilon'^3 U_\xi^\dagger \beta \beta^\dagger \xi^{-1} \mu^\dagger U_\delta & \epsilon'^2 U_\xi^\dagger \beta \beta^\dagger U_\xi \end{pmatrix} \Lambda^{-1} \tag{C.32}$$

$$= \frac{\not{\partial} a}{\langle \phi \rangle} \Lambda^{-1} \frac{\langle \Sigma \rangle^2}{2} \begin{pmatrix} \epsilon^2 \epsilon'^4 \mathcal{A} \mathcal{B} \mathcal{A}^\dagger & \epsilon \epsilon'^3 \mathcal{A} \mathcal{B} \\ \epsilon \epsilon'^3 \mathcal{B} \mathcal{A}^\dagger & \epsilon'^2 \mathcal{B} \end{pmatrix} \Lambda^{-1} \tag{C.33}$$

$$\simeq (\not{\partial} a) \begin{pmatrix} 1/\langle \phi \rangle & v \epsilon' M_q^{-1} \mathcal{A} \mathcal{B} M_F^{-1}/2 \\ v \epsilon' M_F^{-1} \mathcal{B} \mathcal{A}^\dagger M_q^{-1}/2 & \langle \phi \rangle M_F^{-1} \mathcal{B} M_F^{-1}/2 \end{pmatrix}, \tag{C.34}$$

where in the last line we have only kept the leading-order term in ϵ'' as in eq. (C.27), which is of course only justified if $\epsilon'' \ll 1$. So as one would expect, the coupling of SM quarks to axions is only strictly proportional to their masses if the heavy messenger fields are well separated from the PQ scale and can be integrated out. In the final step, we have also used that $\Lambda^{-1} = \text{diag}(M_q^{-1}, M_F^{-1})$. With this, we can finally write the axion couplings as

$$\mathcal{L} \supset \frac{\langle \phi \rangle}{2} (\bar{q}_R \ \bar{F}_R) (\not{\partial} a) \Lambda^{-1} \begin{pmatrix} (\epsilon \epsilon')^2 \mathcal{A} \mathcal{B} \mathcal{A}^\dagger & \epsilon \epsilon' \mathcal{A} \mathcal{B} \\ \epsilon \epsilon' \mathcal{B} \mathcal{A}^\dagger & \mathcal{B} \end{pmatrix} \Lambda^{-1} \begin{pmatrix} q_R \\ F_R \end{pmatrix} \tag{C.35}$$

$$\simeq \frac{1}{\langle \phi \rangle} \bar{q}_R (\not{\partial} a) q_R + \frac{\langle \phi \rangle}{2} \bar{F}_R (M_F^{-1} \mathcal{B} M_F^{-1}) (\not{\partial} a) F_R + \left(\frac{v}{2} \epsilon' \bar{q}_R (M_q^{-1} \mathcal{A} \mathcal{B} M_F^{-1}) (\not{\partial} a) F_R + \text{h.c.} \right), \tag{C.36}$$

in the form that has been used in the main body of this work.

Tree-level contributions to $s \rightarrow d + a$. The previous leading-order expansion only induces flavour-diagonal tree-level couplings between the axion and SM quarks. In this case the $s \rightarrow d + a$ process is only induced at loop level as we have discussed in the main text. However, in general we can not exclude the possibility that higher orders in the expansion of both ϵ and ϵ'' , which induce non-diagonal coupling structures, are relevant. These flavour-violating couplings would trigger the $s \rightarrow d + a$ decay already at tree-level. It therefore becomes important to analyze whether these tree-level couplings can in fact spoil our analysis or, at the very least, whether they restrict the range of values of our expansion parameters.

Before going into the details, let us briefly summarize how large the axial-vector coupling, irrespective of whether it is induced at tree- or loop-level, is allowed to be without being in conflict with the bound $\text{BR}(K^+ \rightarrow \pi^+ + a) < 7.3 \cdot 10^{-11}$ [116]. Generally, a coupling of the form

$$\mathcal{H}_{s \rightarrow da} = \partial_\mu a \bar{d} h_{ds}^S \gamma^\mu (1 + \gamma_5) s + \text{h.c.} \tag{C.37}$$

results in the decay width of eq. (2.7)

$$\Gamma(K^+ \rightarrow \pi^+ a) = \frac{|h_{ds}^S|^2}{16\pi m_{K^+}^3} (m_{K^+}^2 - m_{\pi^+}^2)^2 \lambda^{1/2}(m_{K^+}^2, m_{\pi^+}^2, m_a^2) f_+^2(m_a^2), \quad (\text{C.38})$$

which leads to

$$|h_{ds}^S| \lesssim 1.5 \cdot 10^{-12} \frac{1}{\text{GeV}}. \quad (\text{C.39})$$

From this we can readily conclude that the ϵ expansion is safe: the parameter $\epsilon = v/\langle\phi\rangle$ is $\mathcal{O}(10^{-4})$ for $f_a \sim 10^6$ GeV, which is the region where NA62 is sensitive. The leading-order axial-vector coupling between quarks and the axion in eq. (C.36) therefore corresponds to $1/(2\langle\phi\rangle) \sim 1/(12 \cdot f_a) \sim 1/(10^7 \text{ GeV})$. The NLO contribution to the axial-vector coupling between the axion and SM quarks would be suppressed by an additional factor of ϵ^2 with respect to this leading-order coupling if we were to insert higher orders terms in S in eq. (C.31). This therefore means that the NLO coupling is parametrically suppressed by $10^{-7} \cdot 10^{-8} \text{ GeV}^{-1} = 10^{-15} \text{ GeV}^{-1}$, which is sufficiently far away from the E787 bound quoted in eq. (C.39) and is also out of reach for NA62.

Next, we want to quantify how small ϵ'' has to be in order for significant flavour-violating axion couplings to be avoided at tree level. By expanding ξ^{-1} in ϵ'' as in eq. (C.23), we can write S schematically as

$$S = \begin{pmatrix} \text{unitary} + \mathcal{O}(\epsilon''^2) & \mathcal{O}(\epsilon'') \\ \mathcal{O}(\epsilon'') & \text{unitary} + \mathcal{O}(\epsilon''^2) \end{pmatrix}, \quad (\text{C.40})$$

and estimate the parametric size of the coupling structure in eq. (C.29) as

$$S^\dagger \begin{pmatrix} \mathbf{1} & 0 \\ 0 & 0 \end{pmatrix} S = \begin{pmatrix} \mathbf{1} + \mathcal{O}(\epsilon''^2) & \mathcal{O}(\epsilon'') \\ \mathcal{O}(\epsilon'') & \mathcal{O}(\epsilon''^2) \end{pmatrix}. \quad (\text{C.41})$$

Hence, we can put an upper bound on the tree-level flavour violating axion couplings to the SM quarks,

$$|h_{ds}^{S, \text{treelevel}}| = \frac{1}{2\langle\phi\rangle} \left[S^\dagger \begin{pmatrix} \mathbf{1} & 0 \\ 0 & 0 \end{pmatrix} S \right]_{ds} \lesssim \frac{\epsilon''^2}{2\langle\phi\rangle}. \quad (\text{C.42})$$

To avoid the constraint in (C.39), it is sufficient to require $\epsilon'' \lesssim 10^{-3}$. By further restricting to $\epsilon'' \lesssim 10^{-4}$, the tree-level effect becomes negligible compared to the loop-induced effect. This is exactly the setup which is considered in the main text.

Two more comments are in order. First, note that flavour-violating axion couplings at tree-level are not necessarily a problem but could also be considered an interesting feature of our model. However, since we initially set out to UV complete the model in (4.1) and not to build an axion-flavour model, we choose to restrict ourselves to the case of large scale separations, where the tree-level flavour violation is negligible and flavour violation is induced only by loop processes.

Second, tree-level flavour violation can also be suppressed by the coupling matrices without requiring the masses of the additional fields to be much larger than the PQ scale.

For instance, when $\epsilon' = 0.2$, we can set $\beta = \lambda = \mathbb{1}$ and $\alpha = Y\sqrt{\frac{1+0.2^2}{0.2^2}}$, with Y being the SM Yukawa couplings. This results in vanishing flavour off-diagonal axion couplings to SM quarks in eq. (C.33) at all orders in ϵ'' , as the inversion of ξ in (C.23) is trivial in this case and thus $\mathcal{A}\mathcal{B}\mathcal{A}^\dagger$ becomes proportional to M_q^2 , which is exactly diagonal, without expanding in ϵ'' :

$$M_q^2 = \left[U_\delta^\dagger (\delta - \mu \xi^{-1} \mu^\dagger) U_\delta (\epsilon \epsilon')^2 + \mathcal{O}(\epsilon^3) \right] \frac{\langle \Sigma \rangle^2}{2} \tag{C.43}$$

$$= \left[U_\delta^\dagger \left(\alpha \alpha^\dagger - \frac{1}{1 + \epsilon'^2} \alpha \alpha^\dagger \right) U_\delta (\epsilon \epsilon')^2 + \mathcal{O}(\epsilon^3) \right] \frac{\langle \Sigma \rangle^2}{2} \tag{C.44}$$

$$= \frac{1}{1 + \epsilon'^2} \mathcal{A}\mathcal{B}\mathcal{A}^\dagger \epsilon^2 \epsilon'^4 \frac{\langle \Sigma \rangle^2}{2} + \mathcal{O}(\epsilon^3). \tag{C.45}$$

Moreover, we also see that corrections to the relation in eq. (C.27), crucial to obtain the effective coefficient eq. (5.5), are of the order $\epsilon''^2 = \epsilon'^2 = 4\%$ such that our discussion of the $K^+ \rightarrow \pi^+ + a$ decay is still valid. Note however that α in this explicit realization is already very close to the perturbativity limit of Yukawa couplings.

Electroweak interactions. Right-handed q - and F -quarks are in identical representations of the SM gauge group (considering up- and down-type separately). Hence, the interactions of this chiral components with gauge bosons are unchanged under our transformation. This is not the case for left-handed fields. It is useful to write the 6x6 matrix U as a block matrix

$$U = \begin{pmatrix} A & B \\ C & D \end{pmatrix}, \tag{C.46}$$

where each block is a 3x3 matrix. Note, however, that unitarity of U does not imply unitarity of any of the blocks. Because the transformation above mixes different representations of $SU(2) \times U(1)_Y / U(1)_{EM}$, we have to check how the interactions with W and Z bosons are modified. Let us start with the Z bosons,

$$\begin{aligned} & \begin{pmatrix} \bar{q} \\ \bar{F} \end{pmatrix}^T \gamma_\mu Z^\mu \frac{-g}{\cos(\theta_W)} \left[\pm \frac{1}{2} \begin{pmatrix} \mathbb{1} & 0 \\ 0 & 0 \end{pmatrix} P_L - Q \sin^2(\theta_W) \begin{pmatrix} \mathbb{1} & 0 \\ 0 & \mathbb{1} \end{pmatrix} \right] \begin{pmatrix} q \\ F \end{pmatrix} \\ \rightarrow & \begin{pmatrix} \bar{q} \\ \bar{F} \end{pmatrix}^T \gamma_\mu Z^\mu \frac{-g}{\cos(\theta_W)} \left[\pm \frac{1}{2} \begin{pmatrix} A^\dagger A & A^\dagger B \\ B^\dagger A & B^\dagger B \end{pmatrix} P_L - Q \sin^2(\theta_W) \begin{pmatrix} \mathbb{1} & 0 \\ 0 & \mathbb{1} \end{pmatrix} \right] \begin{pmatrix} q \\ F \end{pmatrix}. \end{aligned} \tag{C.47}$$

In this expression, the upper (lower) sign refers to up- (down-) type quarks and Q is the electromagnetic charge. We see that Z can in principle couple to all available neutral currents, including ones involving light SM quarks of different flavour because A does not have to be unitary. By identifying the blocks A and B in our perturbative result of U in (C.16), we find the Z -interactions at leading order in ϵ to be

$$\mathcal{L} \supset \begin{pmatrix} \bar{q} \\ \bar{F} \end{pmatrix}^T \gamma_\mu Z^\mu \frac{-g}{\cos(\theta_W)} \left[\pm \frac{1}{2} \begin{pmatrix} \mathbb{1} & -\epsilon \epsilon' \mathcal{A} \\ -\epsilon \epsilon' \mathcal{A}^\dagger & (\epsilon \epsilon')^2 \mathcal{A}^\dagger \mathcal{A} \end{pmatrix} P_L - Q \sin^2(\theta_W) \begin{pmatrix} \mathbb{1} & 0 \\ 0 & \mathbb{1} \end{pmatrix} \right] \begin{pmatrix} q \\ F \end{pmatrix}. \tag{C.48}$$

Tree-level flavour-changing couplings to SM quarks only appear at order ϵ^2 .

In the next step, we look at the W interactions. Because these mix up- and down-type quarks, we reintroduce the corresponding labels to write

$$\begin{aligned}
 & \frac{-g}{\sqrt{2}} \begin{pmatrix} \bar{u}_L \\ \bar{d}_L \\ \bar{F}_L^u \\ \bar{F}_L^d \end{pmatrix}^T \gamma^\mu \left[W_\mu^+ \begin{pmatrix} 0 & \mathbb{1} & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} + W_\mu^- \begin{pmatrix} 0 & 0 & 0 & 0 \\ \mathbb{1} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \right] \begin{pmatrix} u_L \\ d_L \\ F_L^u \\ F_L^d \end{pmatrix} \\
 \rightarrow & \frac{-g}{\sqrt{2}} \begin{pmatrix} \bar{u}_L \\ \bar{d}_L \\ \bar{F}_L^u \\ \bar{F}_L^d \end{pmatrix}^T \gamma^\mu \left[W_\mu^+ \begin{pmatrix} 0 & A^{u\dagger} A^d & 0 & A^{u\dagger} B^d \\ 0 & 0 & 0 & 0 \\ 0 & B^{u\dagger} A^d & 0 & B^{u\dagger} B^d \\ 0 & 0 & 0 & 0 \end{pmatrix} + W_\mu^- \begin{pmatrix} 0 & 0 & 0 & 0 \\ A^{d\dagger} A^u & 0 & A^{d\dagger} B^u & 0 \\ 0 & 0 & 0 & 0 \\ B^{d\dagger} A^u & 0 & B^{d\dagger} B^u & 0 \end{pmatrix} \right] \begin{pmatrix} u_L \\ d_L \\ F_L^u \\ F_L^d \end{pmatrix}.
 \end{aligned} \tag{C.49}$$

The W boson couples to all available charged currents. We identify the CKM matrix V as

$$V = A^{u\dagger} A^d = U_\delta^{u\dagger} U_\delta^d + \mathcal{O}(\epsilon^2), \tag{C.50}$$

which unlike in the SM does not have to be unitary, but non-unitarity only appears at order ϵ^2 . When we again insert the perturbative results for A and B , we arrive at

$$\begin{aligned}
 \mathcal{L} \supset & \frac{-g}{\sqrt{2}} \begin{pmatrix} \bar{u}_L \\ \bar{d}_L \\ \bar{F}_L^u \\ \bar{F}_L^d \end{pmatrix}^T \gamma^\mu \left[W_\mu^+ \begin{pmatrix} 0 & V & 0 & -\epsilon\epsilon' V \mathcal{A}_d \\ 0 & 0 & 0 & 0 \\ 0 & -\epsilon\epsilon' \mathcal{A}_u^\dagger V & 0 & (\epsilon\epsilon')^2 \mathcal{A}_u^\dagger V \mathcal{A}_d \\ 0 & 0 & 0 & 0 \end{pmatrix} \right. \\
 & \left. + W_\mu^- \begin{pmatrix} 0 & 0 & 0 & 0 \\ V^\dagger & 0 & -\epsilon\epsilon' V^\dagger \mathcal{A}_u & 0 \\ 0 & 0 & 0 & 0 \\ -\epsilon\epsilon' \mathcal{A}_d^\dagger V^\dagger & 0 & (\epsilon\epsilon')^2 \mathcal{A}_d^\dagger V^\dagger \mathcal{A}_u & 0 \end{pmatrix} \right] \begin{pmatrix} u_L \\ d_L \\ F_L^u \\ F_L^d \end{pmatrix}.
 \end{aligned} \tag{C.51}$$

Radial modes of H and Φ . So far we have not included the radial modes of the Higgs field H and of the PQ field Φ in our discussion. To capture the impact of the diagonalization procedure on their couplings, it is convenient to write

$$\mathcal{L} \supset -(\bar{q}_L \ \bar{F}_L) M_{\text{rad}} \begin{pmatrix} q_R \\ F_R \end{pmatrix} + \text{h.c.}, \tag{C.52}$$

where

$$M_{\text{rad}} = \begin{pmatrix} 0 & \alpha \frac{h}{\sqrt{2}} \\ \beta \frac{\phi}{\sqrt{2}} & 0 \end{pmatrix}. \tag{C.53}$$

The unitary transformation then results in

$$M_{\text{rad}} \rightarrow U^\dagger M_{\text{rad}} S. \tag{C.54}$$

As the expressions are lengthy, we quote each 3x3 blocks separatel

$$\begin{aligned}
 [U^\dagger M_{\text{rad}} S]_{qq} &= \frac{\langle \Sigma \rangle}{2} \left((\epsilon^2 \epsilon'^3 \mathcal{A} \mathcal{B} \mathcal{A}^\dagger \phi + \mathcal{O}(\epsilon^4)) - (\epsilon \epsilon' U_\delta^\dagger \alpha (\lambda^\dagger \xi^{-1} \mu^\dagger - \alpha^\dagger) U_\delta h + \mathcal{O}(\epsilon^3)) \right) M_q^{-1} \\
 &= \frac{\langle \Sigma \rangle}{2} \left(\epsilon^2 \epsilon'^3 \mathcal{A} \mathcal{B} \mathcal{A}^\dagger \phi + \frac{2}{\epsilon \epsilon' \langle \Sigma \rangle^2} M_q^2 h \right) M_q^{-1} \tag{C.55}
 \end{aligned}$$

$$\simeq \frac{M_q}{\langle \phi \rangle} \phi + \frac{M_q}{v} h, \tag{C.56}$$

$$[U^\dagger M_{\text{rad}} S]_{qF} = \frac{\langle \Sigma \rangle}{2} \left((\epsilon \epsilon'^2 \mathcal{A} \mathcal{B} \phi + \mathcal{O}(\epsilon^3)) - (\mathcal{C} h + \mathcal{O}(\epsilon^2)) \right) M_F^{-1}, \tag{C.57}$$

$$\begin{aligned}
 [U^\dagger M_{\text{rad}} S]_{Fq} &= \frac{\langle \Sigma \rangle}{2} \left((\epsilon \epsilon'^2 \mathcal{B} \mathcal{A}^\dagger \phi + \mathcal{O}(\epsilon^3)) + (\epsilon^2 \epsilon'^2 \mathcal{A}^\dagger U_\delta^\dagger \alpha (\lambda^\dagger \xi^{-1} \mu^\dagger - \alpha^\dagger) U_\delta h + \mathcal{O}(\epsilon^3)) \right) M_q^{-1} \\
 &= \frac{\langle \Sigma \rangle}{2} \left(\epsilon \epsilon'^2 \mathcal{B} \mathcal{A}^\dagger \phi - \frac{2}{\langle \Sigma \rangle^2} \mathcal{A}^\dagger M_q^2 h \right) M_q^{-1}, \tag{C.58}
 \end{aligned}$$

$$[U^\dagger M_{\text{rad}} S]_{FF} = \frac{\langle \Sigma \rangle}{2} \left((\epsilon' \mathcal{B} \phi + \mathcal{O}(\epsilon^2)) + (\epsilon \epsilon' \mathcal{A}^\dagger \mathcal{C} h + \mathcal{O}(\epsilon^3)) \right) M_F^{-1}, \tag{C.59}$$

where we have defined $\mathcal{C} = U_\delta^\dagger \alpha \lambda^\dagger U_\xi \simeq 2 \mathcal{A} M_F^2 / \langle \Sigma \rangle^2$. As we can see, modifications of the Higgs coupling to SM quarks always appear together with additional factors of ϵ^2 and are thus negligible. A contribution of ϕ to $s \rightarrow d + a$, on the other hand, must involve internal F -quarks and is therefore proportional to ϵ''^2 . Consequently, any contributions from ϕ loops, in particular the ones related to the unsuppressed coupling in eq. (C.58), become negligible in the $\epsilon'' \ll 1$ limit (as long as this suppression also compensates for possible enhancements originating from the hierarchical coupling matrices). But even if ϵ'' is not small, the ϕ -mediated flavor violation can be suppressed given an explicit form for the α and β coupling matrices. For instance, in the realisation proposed above eq. (C.43), the ϕ -induced loop processes become either flavour-diagonal (because the amplitude has the same flavour structure as the mass matrix) or numerically negligible. This is also true for Higgs-induced loops.

D Counterterm contribution to the kaon decay width

As mentioned in the main text, counterterm contributions are relevant for the computation of the $s \rightarrow d + a$ transition as a result of the renormalization of quark fields (see the right diagram in figure 3). Most important in this context is the renormalization of SM quark fields, while contributions related to heavy F -fields are of higher order in ϵ'' . Note that in the latter case one also encounters divergent loop integrals, which are parametrically suppressed by $\mathcal{O}(\epsilon''^2)$. Cancelling all of these divergences (and, more generally, all one-loop divergences in our UV model) requires a more complete renormalization discussion. However, in the following, we focus on the leading order processes that are not suppressed by ϵ''^2 , and therefore avoid this complication.




Figure 6. Relevant diagrams for determining the renormalization constants contributing to $s \rightarrow d$. The cross marks the insertion of a counterterm. The dots stand for higher order diagrams in our UV model, such as Z boson-induced flavour changes or diagrams with heavy F quarks.

We start by parameterizing the renormalized down-type quark fields as [145]

$$q_L^b = \begin{pmatrix} d \\ s \\ b \end{pmatrix}_L^{\text{bare}} = \left(1 + \frac{1}{2}\delta Z^L\right) \underbrace{\begin{pmatrix} d \\ s \\ b \end{pmatrix}_L}_{q_L}, \quad q_R^b = \begin{pmatrix} d \\ s \\ b \end{pmatrix}_R^{\text{bare}} = \left(1 + \frac{1}{2}\delta Z^R\right) \underbrace{\begin{pmatrix} d \\ s \\ b \end{pmatrix}_R}_{q_R}. \quad (\text{D.1})$$

The renormalization constants of relevance for the $s \rightarrow d + a$ transition are determined by demanding that the one-loop W^\pm contribution to $s \rightarrow d$ and the counterterm cancel each other [127], as depicted in figure 6. The renormalization constants are therefore $\mathcal{O}(g^2)$ as a result of canceling the W^\pm loop.

Inserting the renormalized quark fields from eq. (D.1) into the right-handed derivative axion-fermion interactions, we obtain the terms

$$\mathcal{L} = \frac{1}{\langle\phi\rangle} \bar{q}_R (\not{\partial} a) q_R + \frac{1}{2\langle\phi\rangle} \bar{d}_R \delta Z_{sd}^{R*} (\not{\partial} a) s_R + \frac{1}{2\langle\phi\rangle} \bar{d}_R (\not{\partial} a) \delta Z_{ds}^R s_R + \text{h.c.}, \quad (\text{D.2})$$

where the quark fields now denote the renormalized ones. These new terms involving the renormalization constants therefore induce $s \rightarrow d + a$ at tree-level. Note that this process is of order $g^2/\langle\phi\rangle$, just like the one-loop FCNCs in eq. (5.5). Hence, there is no reason to neglect this counterterm contribution at this stage. As is known in literature, the explicit renormalization calculation does not need to be performed. For this, as will be proven below, the diagrammatic equation in figure 7 can be shown to hold without knowing the precise expressions for the renormalization constants. Moreover, from the renormalization condition shown in figure 6, we can also conclude the two relations showcased in figure 8.

Combining all the relations above, we end up with the key relation depicted in figure 9, which allows us to exchange the counterterm calculation with the computation of additional self-energy diagrams where the ALP is emitted from the external down-type quark legs [146], i.e. the diagrams on the right-hand side. From this point it is most straightforward to simply compute the additional loop diagrams instead of taking a detour to compute renormalization constants. Upon performing the calculations, one notices that the diagrams on the right-hand side of figure 9 add up to zero at linear order in the down- and strange-quark masses. Therefore, the counterterm contribution is subdominant and can safely be neglected for all practical purposes.




Figure 7. Sum of counterterm-induced diagrams for $s \rightarrow d + a$ adding up to zero.




Figure 8. Relation between counterterm insertions and W^\pm loops for $s \rightarrow d + a$.




Figure 9. Relation between $s \rightarrow d + a$ counterterm contribution and ALP emission from external legs.

Proof of the relation in figure 7. The derivative axion-quarks interaction can be generally written as

$$\mathcal{L} \supset \partial_\mu a \bar{q}_R^b G_Q \gamma^\mu q_R^b, \quad (\text{D.3})$$

where $q^b = (d^b, s^b, b^b)^T$ denote the bare down-type quark fields and G_Q is a general coupling matrix of diagonal structure, i.e. $G_Q = \text{diag}(g_d, g_s, g_b)$. Inserting now the renormalized quark fields from eq. (D.1), we obtain

$$\mathcal{L} \supset \partial_\mu a \bar{q}_R \left(1 + \frac{1}{2} \delta Z^{R\dagger}\right) G_Q \left(1 + \frac{1}{2} \delta Z^R\right) \gamma^\mu q_R \quad (\text{D.4})$$

$$= \partial_\mu a \bar{q}_R G_Q \gamma^\mu q_R + \partial_\mu a \bar{q}_R \frac{1}{2} \delta Z^{R\dagger} G_Q \gamma^\mu q_R + \partial_\mu a \bar{q}_R G_Q \frac{1}{2} \delta Z^R \gamma^\mu q_R. \quad (\text{D.5})$$

We also need to trace the influence on the kinetic and mass terms,

$$\mathcal{L} \supset i \bar{q}^b \not{\partial} q^b - \bar{q}^b M q^b \quad (\text{D.6})$$

$$= i \bar{q}_R^b \not{\partial} q_R^b + i \bar{q}_L^b \not{\partial} q_L^b - \bar{q}_R^b M q_L^b - \bar{q}_L^b M q_R^b \quad (\text{D.7})$$

$$= i \bar{q}_R \not{\partial} q_R + i \bar{q}_L \not{\partial} q_L - \bar{q}_R M q_L - \bar{q}_L M q_R + \mathcal{L}_C, \quad (\text{D.8})$$

where $M = \text{diag}(m_d, m_s, m_b)$ and the counterterm Lagrangian \mathcal{L}_C is given by

$$\begin{aligned} \mathcal{L}_C = & \frac{i}{2} \bar{q}_R (\delta Z^R)^\dagger \not{\phi}_{qR} + \frac{i}{2} \bar{q}_R \not{\phi} \delta Z^R q_R + \frac{i}{2} \bar{q}_L (\delta Z^L)^\dagger \not{\phi}_{qL} + \frac{i}{2} \bar{q}_L \not{\phi} \delta Z^L q_L \\ & - \frac{1}{2} \bar{q}_R (\delta Z^R)^\dagger M q_L - \frac{1}{2} \bar{q}_R M \delta Z^L q_L - \frac{1}{2} \bar{q}_L (\delta Z^L)^\dagger M q_R - \frac{1}{2} \bar{q}_L M \delta Z^R q_R. \end{aligned} \quad (\text{D.9})$$

For the left diagram in figure 7, we use eq. (D.4) to obtain

$$i\mathcal{M}_1 = \frac{i(p_1 - p_2)_\mu}{2} \bar{u}_d(p_2) \left(i g_s \delta Z_{sd}^{R*} + i g_d \delta Z_{ds}^R \right) \gamma^\mu P_R u_s(p_1), \quad (\text{D.10})$$

where Z_{sd}^{R*} is a specific matrix element from $\delta Z^{R\dagger}$ and * denotes complex conjugation.

For the diagram in the middle, we can use eqs. (D.4) and (D.9) to find

$$\begin{aligned} i\mathcal{M}_2 = & \bar{u}_d(p_2) \frac{i}{2} \left\{ \left[(\delta Z_{sd}^{R*} + \delta Z_{ds}^R) \not{\phi}_2 - (m_d \delta Z_{ds}^R + m_s \delta Z_{sd}^{L*}) \right] P_R \right. \\ & \left. + \left[(\delta Z_{sd}^{L*} + \delta Z_{ds}^L) \not{\phi}_2 - (m_d \delta Z_{ds}^L + m_s \delta Z_{sd}^{R*}) \right] P_L \right\} \\ & \times \frac{i(\not{\phi}_2 + m_s)}{m_d^2 - m_s^2} i(p_1 - p_2)_\mu i g_s \gamma^\mu P_R u_s(p_1) \end{aligned} \quad (\text{D.11})$$

$$\begin{aligned} = & \frac{-i(p_1 - p_2)_\mu}{2} i g_s \bar{u}_d(p_2) \left\{ \left[\delta Z_{sd}^{R*} m_d - m_s \delta Z_{sd}^{L*} \right] P_R + \left[\delta Z_{sd}^{L*} m_d - m_s \delta Z_{sd}^{R*} \right] P_L \right\} \\ & \times \frac{\not{\phi}_2 + m_s}{m_d^2 - m_s^2} \gamma^\mu P_R u_s(p_1), \end{aligned} \quad (\text{D.12})$$

where we have used the equation of motion $\bar{u}_d(p_2) \not{\phi}_2 = \bar{u}_d(p_2) m_d$. Using the equation of motion once more for $\not{\phi}_2$ in the propagator yields

$$\begin{aligned} i\mathcal{M}_2 = & \frac{-i(p_1 - p_2)_\mu}{2} i g_s \bar{u}_d(p_2) \left\{ \left(\left[\delta Z_{sd}^{R*} m_d - m_s \delta Z_{sd}^{L*} \right] P_R + \left[\delta Z_{sd}^{L*} m_d - m_s \delta Z_{sd}^{R*} \right] P_L \right) \frac{m_s}{m_d^2 - m_s^2} \right. \\ & \left. + \left(\left[\delta Z_{sd}^{R*} m_d - m_s \delta Z_{sd}^{L*} \right] P_L + \left[\delta Z_{sd}^{L*} m_d - m_s \delta Z_{sd}^{R*} \right] P_R \right) \frac{m_d}{m_d^2 - m_s^2} \right\} \gamma^\mu P_R u_s(p_1) \end{aligned} \quad (\text{D.13})$$

$$\begin{aligned} = & \frac{-i(p_1 - p_2)_\mu}{2} i g_s \bar{u}_d(p_2) \left\{ -m_s^2 \delta Z_{sd}^{L*} P_R - m_s^2 \delta Z_{sd}^{R*} P_L \right. \\ & \left. + \delta Z_{sd}^{R*} m_d^2 P_L + \delta Z_{sd}^{L*} m_d^2 P_R \right\} \frac{1}{m_d^2 - m_s^2} \gamma^\mu P_R u_s(p_1) \end{aligned} \quad (\text{D.14})$$

$$= \frac{-i(p_1 - p_2)_\mu}{2} i g_s \bar{u}_d(p_2) \delta Z_{sd}^{R*} \gamma^\mu P_R u_s(p_1), \quad (\text{D.15})$$

which cancels the g_s term in eq. (D.10). Note that the terms involving δZ^L cancel due to $P_R P_L = 0$.




Figure 10. One-loop contributions to photon interactions with quarks in our UV model, which are of relevance for the EDM discussion.

We can proceed in a similar fashion for the right diagram in figure 7,

$$\begin{aligned}
 i\mathcal{M}_3 &= \bar{u}_d(p_2) i(p_1 - p_2)_\mu i g_d \gamma^\mu P_R \frac{i(\not{p}_1 + m_d)}{m_s^2 - m_d^2} \frac{i}{2} \\
 &\times \left\{ \left[(\delta Z_{sd}^{R*} + \delta Z_{ds}^R) \not{p}_1 - (m_d \delta Z_{ds}^R + m_s \delta Z_{sd}^{L*}) \right] P_R \right. \\
 &\quad \left. + \left[(\delta Z_{sd}^{L*} + \delta Z_{ds}^L) \not{p}_1 - (m_d \delta Z_{ds}^L + m_s \delta Z_{sd}^{R*}) \right] P_L \right\} u_s(p_1) \quad (\text{D.16})
 \end{aligned}$$

$$\begin{aligned}
 &= \frac{-i(p_1 - p_2)_\mu}{2} i g_d \bar{u}_d(p_2) \gamma^\mu P_R \frac{\not{p}_1 + m_d}{m_s^2 - m_d^2} \\
 &\times \left\{ \delta Z_{ds}^R m_s P_L - m_d \delta Z_{ds}^R P_R + \delta Z_{ds}^L m_s P_R - m_d \delta Z_{ds}^L P_L \right\} u_s(p_1) \quad (\text{D.17})
 \end{aligned}$$

$$\begin{aligned}
 &= \frac{-i(p_1 - p_2)_\mu}{2} i g_d \bar{u}_d(p_2) \gamma^\mu P_R \\
 &\times \left\{ \frac{m_d}{m_s^2 - m_d^2} \left(\delta Z_{ds}^R m_s P_L - m_d \delta Z_{ds}^R P_R + \delta Z_{ds}^L m_s P_R - m_d \delta Z_{ds}^L P_L \right) \right. \\
 &\quad \left. + \frac{m_s}{m_s^2 - m_d^2} \left(\delta Z_{ds}^R m_s P_R - m_d \delta Z_{ds}^R P_L + \delta Z_{ds}^L m_s P_L - m_d \delta Z_{ds}^L P_R \right) \right\} u_s(p_1) \quad (\text{D.18})
 \end{aligned}$$

$$= \frac{-i(p_1 - p_2)_\mu}{2} i g_d \bar{u}_d(p_2) \gamma^\mu \delta Z_{ds}^R P_R u_s(p_1), \quad (\text{D.19})$$

which cancels the g_d term in eq. (D.10). We therefore conclude that the sum of all three diagrams vanishes,

$$i\mathcal{M}_1 + i\mathcal{M}_2 + i\mathcal{M}_3 = 0, \quad (\text{D.20})$$

independently of the precise expression of the renormalization constants.

E EDM calculations

General considerations. In this appendix, we assess whether the one-loop contributions shown in figures 10 and 11 can generate quark EDMs and hence also a neutron EDM. Before computing the relevant diagrams explicitly, it is worthwhile to remember the discussion of appendix A. Similar to the case of weak interactions, CP-violating interactions can be parameterized by assigning complex phases to the corresponding operators. Therefore, once a specific operator is inserted twice in a given diagram, i.e. a specific vertex and its hermitian conjugate appear together, the CP phase necessarily drops out.




Figure 11. Barr-Zee diagrams relevant for the quark EDM computation in our UV model.

To be more specific, consider the diagram on the left in figure 10. Let us first write the corresponding $a - q - F$ vertices generally as $(\partial_\mu a) \bar{q}_i \gamma^\mu X_{ij} P_R F_j + \text{h.c.}$ with an arbitrary 3×3 flavour matrix X , and as usual use q (F) to denote the up- or down-type quark triplets in the SM (F -quark) sector. Note that the right-handed coupling structure is fixed by our diagonalization procedure, see appendix C. This operator violates CP if X has complex-valued entries. Note now that the external quarks are fixed in the Feynman diagram and hence the same $a - q - F$ vertex appears twice in the one-loop diagram because the photon only couples flavour-diagonally to all orders in ϵ and ϵ'' . If we rewrite the relevant entry X_{ql} as $x_{ql} e^{i\theta_{ql}}$, where l denotes the internal F -quark flavour and θ_{ql} denotes the CP violating phase, we get $e^{i\theta_{ql}} \cdot e^{-i\theta_{ql}} = 1$ in the Feynman amplitude and hence the CP violating phase drops out. As a result, the diagram on the left in figure 10 cannot produce any net CP-violating quark EDM. This is also consistent with a description based on Jarlskog invariants, discussed in ref. [129], upon taking the right-handed coupling structure into account.

An analogous argument can be applied to the remaining diagrams in figure 10, and also to the ones where the photon is radiated off an external quark line in those diagrams. None of these diagrams can therefore induce an EDM at any order in ϵ and ϵ'' . This is true even if further CP-violating coupling structures were to occur at higher orders, which could ‘allow’ for an internal SM quark $q' \neq q$ to appear in the loops (if $q' = q$, no CP violating vertex is possible as discussed in appendix A).

Thus, only the Barr-Zee diagram on the left of figure 11 and its ‘conjugate’ with the internal gauge boson and axion line interchanged remain.¹⁵ For these diagrams the reasoning is not as straightforward since no vertex appears twice in the diagram. For an internal photon line, however, one notices that no CP violating vertex appears in the diagram. As the photon only couples flavour-diagonally, the internal quark line has to be the same quark q as the external one. In this case, the axion-quark vertex cannot violate CP (see appendix A) and consequently this diagram is fully CP conserving. This reasoning does not apply to the Barr-Zee diagram with an internal Z -boson and hence a contribution to the quark EDM d_q given by $-(d_q/2) F^{\mu\nu} \bar{q} \sigma_{\mu\nu} i \gamma_5 q$ can be expected. For the explicit computation below, we follow the momentum flow as shown in the middle and right panels of figure 11.

Z-induced Barr-Zee diagram. We now calculate the EDM contribution based on the Barr-Zee diagram with an internal Z boson in figure 11. Considering the order in ϵ and ϵ'' to which we have expanded, to obtain a net CP violating contribution we have to focus on an internal F -quark for the fermionic line. We parameterize the flavour matrix structure of

¹⁵Formally, by counting the coupling constant insertions, these diagrams are of the same order as other two-loop diagrams. We still consider them explicitly as only a single loop integration is involved.

the $a - q - F$ and $Z - q - F$ vertices by $X_{ql} = \epsilon' v (M_q^{-1} \mathcal{A} \mathcal{B} M_F^{-1})_{ql} / 2$ and $Y_{ql} = \frac{g}{2c_W} \epsilon \epsilon' \mathcal{A}_{ql}$, respectively, as dictated by eqs. (4.22) and (5.2). For both diagrams and for each individual F -quark with index l , the amplitude reads

$$\begin{aligned}
 i\mathcal{M} \simeq & (\pm) \int \frac{d^4 k}{(2\pi)^4} \left[\bar{u}_q(p_2) i X_{ql} i (\not{k} + \not{p}_1 - \not{p}_2) P_R \frac{i(\not{k} + \not{p}_1 + m_{F_l})}{(k + p_1)^2 - m_{F_l}^2} i \gamma_\mu Y_{lq}^\dagger P_L u_q(p_1) \right] \\
 & \times \frac{-i \left(g^{\mu\nu} - \frac{k^\mu k^\nu}{m_Z^2} \right)}{k^2 - m_Z^2} \frac{e^2}{(4\pi)^2 f_a} i \epsilon^{\beta\nu\lambda\alpha} (i k_\beta) (i(p_1 - p_2)_\lambda) \epsilon_\alpha^* \frac{i}{(k + p_1 - p_2)^2 - m_a^2} \\
 & (\pm) \int \frac{d^4 k}{(2\pi)^4} \left[\bar{u}_q(p_2) i \gamma_\mu Y_{ql} P_L \frac{i(\not{k} + \not{p}_1 + m_{F_l})}{(k + p_1)^2 - m_{F_l}^2} i X_{lq}^\dagger (-i \not{k}) P_R u_q(p_1) \right] \\
 & \times \frac{-i \left(g^{\mu\nu} - \frac{(k+p_1-p_2)^\mu (k+p_1-p_2)^\nu}{m_Z^2} \right)}{(k + p_1 - p_2)^2 - m_Z^2} \frac{i}{k^2 - m_a^2} \frac{e^2}{(4\pi)^2 f_a} \\
 & \times i \epsilon^{\beta\nu\lambda\alpha} (-i) (k + p_1 - p_2)_\beta (i(p_1 - p_2)_\lambda) \epsilon_\alpha^*. \tag{E.1}
 \end{aligned}$$

Here, we have schematically used the Feynman rule due to $(e/4\pi)^2 (a/f_a) F^{\mu\nu} \tilde{Z}_{\mu\nu}$ resulting from the chiral rotation discussed in section 4.3. The global \pm depends on whether q is an up- or down-type quark, see eq. (5.2), and ϵ_α^* is the photon polarization vector. We can write the amplitude in a more compact form as

$$\begin{aligned}
 i\mathcal{M} \simeq & (\pm) (-i) \frac{e^2}{(4\pi)^2 f_a} \int \frac{d^4 k}{(2\pi)^4} X_{ql} Y_{lq}^\dagger \epsilon^{\beta\nu\lambda\alpha} k_\beta (p_1 - p_2)_\lambda \epsilon_\alpha^* \left[\bar{u}_q(p_2) (\not{k} + \not{p}_1 - \not{p}_2) P_R \right. \\
 & \times \left. \frac{m_{F_l}}{(k + p_1)^2 - m_{F_l}^2} \gamma_\mu u_q(p_1) \right] \frac{g^{\mu\nu} - \frac{k^\mu k^\nu}{m_Z^2}}{k^2 - m_Z^2} \frac{1}{(k + p_1 - p_2)^2 - m_a^2} \\
 & (\pm) (-i) \frac{e^2}{(4\pi)^2 f_a} \int \frac{d^4 k}{(2\pi)^4} Y_{ql} X_{lq}^\dagger \epsilon^{\beta\nu\lambda\alpha} (k + p_1 - p_2)_\beta (p_1 - p_2)_\lambda \epsilon_\alpha^* \left[\bar{u}_q(p_2) \gamma_\mu P_L \right. \\
 & \times \left. \frac{m_{F_l}}{(k + p_1)^2 - m_{F_l}^2} \not{k} u_q(p_1) \right] \frac{g^{\mu\nu} - \frac{(k+p_1-p_2)^\mu (k+p_1-p_2)^\nu}{m_Z^2}}{(k + p_1 - p_2)^2 - m_Z^2} \frac{1}{k^2 - m_a^2}. \tag{E.2}
 \end{aligned}$$

We then use `Package-X` [126] to compute the loop integrals and map out the relevant Lorentz structure for the quark EDM. Both of the Barr-Zee diagrams give the same loop functions $f(m_{F_l}, m_a, m_Z, m_q) \equiv f(m_{F_l}, m_q)$ for the EDM operator, differing only in their sign. The relevant amplitude structure therefore reads

$$i\mathcal{M} \simeq (-i) \frac{e^2}{(4\pi)^2 f_a} (X_{ql} Y_{lq}^\dagger - Y_{ql} X_{lq}^\dagger) \frac{i}{16\pi^2} m_{F_l} f(m_{F_l}, m_q) \bar{u}_q(p_2) \frac{i\sigma^{\alpha\beta}\gamma_5}{2m_q} \epsilon_\alpha^* (p_1 - p_2)_\beta u_q(p_1). \tag{E.3}$$

Inserting the LO expressions for X and Y , we get

$$\begin{aligned}
 i\mathcal{M} \simeq & (-i) \frac{e^2}{(4\pi)^2 f_a} \epsilon \epsilon'^2 \frac{v}{2} \frac{g}{2c_W} \frac{1}{m_q} ((\mathcal{A}\mathcal{B})_{ql} \mathcal{A}_{lq}^\dagger - \mathcal{A}_{ql} (\mathcal{B}\mathcal{A}^\dagger)_{lq}) \\
 & \times \frac{i}{16\pi^2} f(m_{F_l}, m_q) \bar{u}_q(p_2) \frac{i\sigma^{\alpha\beta}\gamma_5}{2m_q} \epsilon_\alpha^* (p_1 - p_2)_\beta u_q(p_1). \tag{E.4}
 \end{aligned}$$

In the case considered in the main text, where all F quarks are assumed to be equally heavy, performing the sum over all internal F -quarks, (i.e. over the index l) leads to a vanishing amplitude. Higher orders in the expansion of the relevant $Z - q - F$ and $a - q - F$ vertices can in principle again lead to a non-zero result, but one that would be suppressed by higher powers of ϵ and would therefore be negligible.

As the Z-induced Barr-Zee diagram is formally of two-loop order and to motivate that the estimate for the upper limit of the neutron EDM in eq. (5.6) is reasonable, we can, for the sake of an estimate, assume that all F quarks have different mass and hence no cancellation takes place. The loop function $f(m_{F_l}, m_q)$ is UV divergent and hence exhibits a scale dependence μ . This dependence would of course disappear if further contributions were taken into account. Expanding in inverse powers of m_{F_l} , we find

$$f(m_{F_l}, m_q) \approx \frac{im_q}{4} \left(-3 + 2 \log \left(\frac{\mu^2}{m_{F_l}^2} \right) \right) + \dots \quad (\text{E.5})$$

Without performing a full two-loop analysis, we conservatively assume the term in the brackets to be an $\mathcal{O}(1)$ factor. The remaining finite terms in $f(m_{F_l})$, which are not expected to fully cancel, would actually be significantly smaller for $m_F \geq 1 \text{ TeV}$. Moreover, if we parameterize the entries of $(\mathcal{AB})_{ql}$ and \mathcal{A}_{ql} by $x_{ql} e^{i\theta_{ql}}$ and $a_{ql} e^{i\phi_{ql}}$, respectively, we obtain

$$i\mathcal{M} \simeq (-i) \frac{e^2}{(4\pi)^2 f_a} \epsilon \epsilon'^2 \frac{v}{2} \frac{g}{2c_W} \frac{1}{m_q} \left(x_{ql} e^{i\theta_{ql}} a_{ql} e^{-i\phi_{ql}} - a_{ql} e^{i\phi_{ql}} x_{ql} e^{-i\theta_{ql}} \right) \times \frac{i}{16\pi^2} \frac{im_q}{4} \bar{u}_q(p_2) \frac{i\sigma^{\alpha\beta}\gamma_5}{2m_q} \epsilon_\alpha^* (p_1 - p_2)_\beta u_q(p_1). \quad (\text{E.6})$$

The EDM operator to map onto is $-(d_q/2) F^{\mu\nu} \bar{q} \sigma_{\mu\nu} i\gamma_5 q$, where d_q has to be real-valued for the operator to be hermitian. By comparing with eq. (E.6) and including the contributions from all three F quarks, we have

$$d_q \sim 3 \cdot \frac{1}{128\pi^2} \frac{e^2}{(4\pi)^2 f_a} \epsilon \epsilon'^2 \frac{v}{2} \frac{g}{2c_W} \frac{1}{m_q} x_{ql} a_{ql} 2 \sin(\theta_{ql} - \phi_{ql}) \quad (\text{E.7})$$

$$\sim 3 \cdot \frac{1}{128\pi^2} \frac{\alpha}{4\pi f_a} \epsilon \epsilon'^2 \frac{v}{2} \frac{e}{2s_W c_W} \frac{1}{m_q} \frac{2m_q^2}{\epsilon^2 \epsilon'^4 \langle \Sigma \rangle^2} 2 \sin(\theta_{ql} - \phi_{ql}) \quad (\text{E.8})$$

$$\sim 3 \cdot \frac{1}{128\pi^2} \frac{\alpha}{4\pi f_a} \frac{1}{6f_a} \frac{e}{2s_W c_W} m_q 2 \sin(\theta_{ql} - \phi_{ql}), \quad (\text{E.9})$$

where we have identified $x_{ql} a_{ql} \sim 2m_q^2 / (\epsilon^2 \epsilon'^4 \langle \Sigma \rangle^2)$ in the second line based on eq. (C.27). For the up- and down-quarks respectively, fixing $f_a = 4 \cdot 10^6 \text{ GeV}$ and assuming $\sin(\theta_{ql} - \phi_{ql}) \approx 1$, we get

$$d_d \sim 3 \cdot 10^{-36} e \cdot \text{cm}, \quad d_u \sim 1 \cdot 10^{-36} e \cdot \text{cm}. \quad (\text{E.10})$$

This result is in agreement with the naive estimate for the upper limit on the neutron EDM in eq. (5.6) for $d_n^{\text{UV}} \sim (4/3)d_d - (1/3)d_u$ [130]. Of course, it should be kept in mind that the result in eq. (E.10) is only an estimate as well. However, given that the sensitivity of current neutron EDM experiments is away by several orders of magnitude, it is an estimate that is sufficient for our purposes.

Open Access. This article is distributed under the terms of the Creative Commons Attribution License ([CC-BY 4.0](https://creativecommons.org/licenses/by/4.0/)), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited.

References

- [1] G. Buchalla, A.J. Buras and M.E. Lautenbacher, *Weak decays beyond leading logarithms*, *Rev. Mod. Phys.* **68** (1996) 1125 [[hep-ph/9512380](#)] [[INSPIRE](#)].
- [2] J. Aebischer, M. Fael, C. Greub and J. Virto, *B physics Beyond the Standard Model at One Loop: Complete Renormalization Group Evolution below the Electroweak Scale*, *JHEP* **09** (2017) 158 [[arXiv:1704.06639](#)] [[INSPIRE](#)].
- [3] E.E. Jenkins, A.V. Manohar and P. Stoffer, *Low-Energy Effective Field Theory below the Electroweak Scale: Operators and Matching*, *JHEP* **03** (2018) 016 [[arXiv:1709.04486](#)] [[INSPIRE](#)].
- [4] W. Buchmüller and D. Wyler, *Effective Lagrangian Analysis of New Interactions and Flavor Conservation*, *Nucl. Phys. B* **268** (1986) 621 [[INSPIRE](#)].
- [5] B. Grzadkowski, M. Iskrzynski, M. Misiak and J. Rosiek, *Dimension-Six Terms in the Standard Model Lagrangian*, *JHEP* **10** (2010) 085 [[arXiv:1008.4884](#)] [[INSPIRE](#)].
- [6] F. Feruglio, *The Chiral approach to the electroweak interactions*, *Int. J. Mod. Phys. A* **8** (1993) 4937 [[hep-ph/9301281](#)] [[INSPIRE](#)].
- [7] R. Alonso, M.B. Gavela, L. Merlo, S. Rigolin and J. Yepes, *The Effective Chiral Lagrangian for a Light Dynamical “Higgs Particle”*, *Phys. Lett. B* **722** (2013) 330 [Erratum *ibid.* **726** (2013) 926] [[arXiv:1212.3305](#)] [[INSPIRE](#)].
- [8] G. Buchalla, O. Catà and C. Krause, *Complete Electroweak Chiral Lagrangian with a Light Higgs at NLO*, *Nucl. Phys. B* **880** (2014) 552 [Erratum *ibid.* **913** (2016) 475] [[arXiv:1307.5017](#)] [[INSPIRE](#)].
- [9] P. Sikivie, *Experimental Tests of the Invisible Axion*, *Phys. Rev. Lett.* **51** (1983) 1415 [Erratum *ibid.* **52** (1984) 695] [[INSPIRE](#)].
- [10] CAST collaboration, *New CAST Limit on the Axion-Photon Interaction*, *Nature Phys.* **13** (2017) 584 [[arXiv:1705.02290](#)] [[INSPIRE](#)].
- [11] IAXO collaboration, *Physics potential of the International Axion Observatory (IAXO)*, *JCAP* **06** (2019) 047 [[arXiv:1904.09155](#)] [[INSPIRE](#)].
- [12] G. Raffelt, *Stars as laboratories for fundamental physics: The astrophysics of neutrinos, axions, and other weakly interacting particles*, University of Chicago Press (1996) [[INSPIRE](#)].
- [13] G.G. Raffelt, *Astrophysical axion bounds*, *Lect. Notes Phys.* **741** (2008) 51 [[hep-ph/0611350](#)] [[INSPIRE](#)].
- [14] A. Arvanitaki and S. Dubovsky, *Exploring the String Axiverse with Precision Black Hole Physics*, *Phys. Rev. D* **83** (2011) 044026 [[arXiv:1004.3558](#)] [[INSPIRE](#)].
- [15] A. Arvanitaki, M. Baryakhtar and X. Huang, *Discovering the QCD Axion with Black Holes and Gravitational Waves*, *Phys. Rev. D* **91** (2015) 084011 [[arXiv:1411.2263](#)] [[INSPIRE](#)].
- [16] J. Preskill, M.B. Wise and F. Wilczek, *Cosmology of the Invisible Axion*, *Phys. Lett. B* **120** (1983) 127 [[INSPIRE](#)].

- [17] L.F. Abbott and P. Sikivie, *A Cosmological Bound on the Invisible Axion*, *Phys. Lett. B* **120** (1983) 133 [[INSPIRE](#)].
- [18] M. Dine and W. Fischler, *The Not So Harmless Axion*, *Phys. Lett. B* **120** (1983) 137 [[INSPIRE](#)].
- [19] ADMX collaboration, *Results from a high sensitivity search for cosmic axions*, *Phys. Rev. Lett.* **80** (1998) 2043 [[astro-ph/9801286](#)] [[INSPIRE](#)].
- [20] D. Horns, J. Jaeckel, A. Lindner, A. Lobanov, J. Redondo and A. Ringwald, *Searching for WISPy Cold Dark Matter with a Dish Antenna*, *JCAP* **04** (2013) 016 [[arXiv:1212.2970](#)] [[INSPIRE](#)].
- [21] D. Budker, P.W. Graham, M. Ledbetter, S. Rajendran and A. Sushkov, *Proposal for a Cosmic Axion Spin Precession Experiment (CASPEr)*, *Phys. Rev. X* **4** (2014) 021030 [[arXiv:1306.6089](#)] [[INSPIRE](#)].
- [22] J. Jaeckel and J. Redondo, *An antenna for directional detection of WISPy dark matter*, *JCAP* **11** (2013) 016 [[arXiv:1307.7181](#)] [[INSPIRE](#)].
- [23] W. Chung, *CULTASK, The Coldest Axion Experiment at CAPP/IBS in Korea*, *PoS CORFU2015* (2016) 047 [[INSPIRE](#)].
- [24] Y. Kahn, B.R. Safdi and J. Thaler, *Broadband and Resonant Approaches to Axion Dark Matter Detection*, *Phys. Rev. Lett.* **117** (2016) 141801 [[arXiv:1602.01086](#)] [[INSPIRE](#)].
- [25] MADMAX WORKING GROUP collaboration, *Dielectric Haloscopes: A New Way to Detect Axion Dark Matter*, *Phys. Rev. Lett.* **118** (2017) 091801 [[arXiv:1611.05865](#)] [[INSPIRE](#)].
- [26] D. Alesini, D. Babusci, D. Di Gioacchino, C. Gatti, G. Lamanna and C. Ligi, *The KLASH Proposal*, [arXiv:1707.06010](#) [[INSPIRE](#)].
- [27] A.A. Melcón et al., *Axion Searches with Microwave Filters: the RADES project*, *JCAP* **05** (2018) 040 [[arXiv:1803.01243](#)] [[INSPIRE](#)].
- [28] ADMX collaboration, *A Search for Invisible Axion Dark Matter with the Axion Dark Matter Experiment*, *Phys. Rev. Lett.* **120** (2018) 151301 [[arXiv:1804.05750](#)] [[INSPIRE](#)].
- [29] ADMX collaboration, *Extended Search for the Invisible Axion with the Axion Dark Matter Experiment*, *Phys. Rev. Lett.* **124** (2020) 101303 [[arXiv:1910.08638](#)] [[INSPIRE](#)].
- [30] I.G. Irastorza and J. Redondo, *New experimental approaches in the search for axion-like particles*, *Prog. Part. Nucl. Phys.* **102** (2018) 89 [[arXiv:1801.08127](#)] [[INSPIRE](#)].
- [31] M.T. Ressell, *Limits to the radiative decay of the axion*, *Phys. Rev. D* **44** (1991) 3001 [[INSPIRE](#)].
- [32] M.A. Bershadsky, M.T. Ressell and M.S. Turner, *Telescope search for multi-eV axions*, *Phys. Rev. Lett.* **66** (1991) 1398 [[INSPIRE](#)].
- [33] J.M. Overduin and P.S. Wesson, *Dark matter and background light*, *Phys. Rept.* **402** (2004) 267 [[astro-ph/0407207](#)] [[INSPIRE](#)].
- [34] D. Grin, G. Covone, J.-P. Kneib, M. Kamionkowski, A. Blain and E. Jullo, *A Telescope Search for Decaying Relic Axions*, *Phys. Rev. D* **75** (2007) 105018 [[astro-ph/0611502](#)] [[INSPIRE](#)].
- [35] A. Boyarsky, A. Neronov, O. Ruchayskiy, M. Shaposhnikov and I. Tkachev, *Where to find a dark matter sterile neutrino?*, *Phys. Rev. Lett.* **97** (2006) 261302 [[astro-ph/0603660](#)] [[INSPIRE](#)].

- [36] A. Boyarsky, O. Ruchayskiy and M. Shaposhnikov, *The Role of sterile neutrinos in cosmology and astrophysics*, *Ann. Rev. Nucl. Part. Sci.* **59** (2009) 191 [[arXiv:0901.0011](#)] [[INSPIRE](#)].
- [37] G. Vertongen and C. Weniger, *Hunting Dark Matter Gamma-Ray Lines with the Fermi LAT*, *JCAP* **05** (2011) 027 [[arXiv:1101.2610](#)] [[INSPIRE](#)].
- [38] D. Cadamuro and J. Redondo, *Cosmological bounds on pseudo Nambu-Goldstone bosons*, *JCAP* **02** (2012) 032 [[arXiv:1110.2895](#)] [[INSPIRE](#)].
- [39] P. Arias, D. Cadamuro, M. Goodsell, J. Jaeckel, J. Redondo and A. Ringwald, *WISPy Cold Dark Matter*, *JCAP* **06** (2012) 013 [[arXiv:1201.5902](#)] [[INSPIRE](#)].
- [40] J. Jaeckel, J. Redondo and A. Ringwald, *3.55 keV hint for decaying axionlike particle dark matter*, *Phys. Rev. D* **89** (2014) 103511 [[arXiv:1402.7335](#)] [[INSPIRE](#)].
- [41] A. Hook, Y. Kahn, B.R. Safdi and Z. Sun, *Radio Signals from Axion Dark Matter Conversion in Neutron Star Magnetospheres*, *Phys. Rev. Lett.* **121** (2018) 241102 [[arXiv:1804.03145](#)] [[INSPIRE](#)].
- [42] A. Caputo, C.P. Garay and S.J. Witte, *Looking for Axion Dark Matter in Dwarf Spheroidals*, *Phys. Rev. D* **98** (2018) 083024 [*Erratum ibid.* **99** (2019) 089901] [[arXiv:1805.08780](#)] [[INSPIRE](#)].
- [43] A. Caputo, M. Regis, M. Taoso and S.J. Witte, *Detecting the Stimulated Decay of Axions at RadioFrequencies*, *JCAP* **03** (2019) 027 [[arXiv:1811.08436](#)] [[INSPIRE](#)].
- [44] J.W. Foster et al., *Green Bank and Effelsberg Radio Telescope Searches for Axion Dark Matter Conversion in Neutron Star Magnetospheres*, *Phys. Rev. Lett.* **125** (2020) 171301 [[arXiv:2004.00011](#)] [[INSPIRE](#)].
- [45] J.-W. Wang, X.-J. Bi, R.-M. Yao and P.-F. Yin, *Exploring axion dark matter through radio signals from magnetic white dwarf stars*, *Phys. Rev. D* **103** (2021) 115021 [[arXiv:2101.02585](#)] [[INSPIRE](#)].
- [46] P. Agrawal and K. Howe, *Factoring the Strong CP Problem*, *JHEP* **12** (2018) 029 [[arXiv:1710.04213](#)] [[INSPIRE](#)].
- [47] D.S.M. Alves and N. Weiner, *A viable QCD axion in the MeV mass range*, *JHEP* **07** (2018) 092 [[arXiv:1710.03764](#)] [[INSPIRE](#)].
- [48] M.K. Gaillard, M.B. Gavela, R. Houtz, P. Quilez and R. Del Rey, *Color unified dynamical axion*, *Eur. Phys. J. C* **78** (2018) 972 [[arXiv:1805.06465](#)] [[INSPIRE](#)].
- [49] T. Gherghetta, V.V. Khoze, A. Pomarol and Y. Shirman, *The Axion Mass from 5D Small Instantons*, *JHEP* **03** (2020) 063 [[arXiv:2001.05610](#)] [[INSPIRE](#)].
- [50] J. Jaeckel, M. Jankowiak and M. Spannowsky, *LHC probes the hidden sector*, *Phys. Dark Univ.* **2** (2013) 111 [[arXiv:1212.3620](#)] [[INSPIRE](#)].
- [51] K. Mimasu and V. Sanz, *ALPs at Colliders*, *JHEP* **06** (2015) 173 [[arXiv:1409.4792](#)] [[INSPIRE](#)].
- [52] J. Jaeckel and M. Spannowsky, *Probing MeV to 90 GeV axion-like particles with LEP and LHC*, *Phys. Lett. B* **753** (2016) 482 [[arXiv:1509.00476](#)] [[INSPIRE](#)].
- [53] I. Brivio et al., *ALPs Effective Field Theory and Collider Signatures*, *Eur. Phys. J. C* **77** (2017) 572 [[arXiv:1701.05379](#)] [[INSPIRE](#)].

- [54] M. Bauer, M. Neubert and A. Thamm, *Collider Probes of Axion-Like Particles*, *JHEP* **12** (2017) 044 [[arXiv:1708.00443](#)] [[INSPIRE](#)].
- [55] S. Knapen, T. Lin, H.K. Lou and T. Melia, *Searching for Axionlike Particles with Ultrapерipheral Heavy-Ion Collisions*, *Phys. Rev. Lett.* **118** (2017) 171801 [[arXiv:1607.06083](#)] [[INSPIRE](#)].
- [56] CMS collaboration, *Evidence for light-by-light scattering and searches for axion-like particles in ultraperipheral PbPb collisions at $\sqrt{s_{NN}} = 5.02$ TeV*, *Phys. Lett. B* **797** (2019) 134826 [[arXiv:1810.04602](#)] [[INSPIRE](#)].
- [57] ATLAS collaboration, *Measurement of light-by-light scattering and search for axion-like particles with 2.2 nb^{-1} of Pb + Pb data with the ATLAS detector*, *JHEP* **03** (2021) 243 [[arXiv:2008.05355](#)] [[INSPIRE](#)].
- [58] A. Mariotti, D. Redigolo, F. Sala and K. Tobioka, *New LHC bound on low-mass diphoton resonances*, *Phys. Lett. B* **783** (2018) 13 [[arXiv:1710.01743](#)] [[INSPIRE](#)].
- [59] M. Freytsis, Z. Ligeti and J. Thaler, *Constraining the Axion Portal with $B \rightarrow Kl^{+l^{-}}$* , *Phys. Rev. D* **81** (2010) 034001 [[arXiv:0911.5355](#)] [[INSPIRE](#)].
- [60] E. Izaguirre, T. Lin and B. Shuve, *Searching for Axionlike Particles in Flavor-Changing Neutral Current Processes*, *Phys. Rev. Lett.* **118** (2017) 111802 [[arXiv:1611.09355](#)] [[INSPIRE](#)].
- [61] M.J. Dolan, T. Ferber, C. Hearty, F. Kahlhoefer and K. Schmidt-Hoberg, *Revised constraints and Belle II sensitivity for visible and invisible axion-like particles*, *JHEP* **12** (2017) 094 [*Erratum ibid.* **03** (2021) 190] [[arXiv:1709.00009](#)] [[INSPIRE](#)].
- [62] X. Cid Vidal, A. Mariotti, D. Redigolo, F. Sala and K. Tobioka, *New Axion Searches at Flavor Factories*, *JHEP* **01** (2019) 113 [*Erratum ibid.* **06** (2020) 141] [[arXiv:1810.09452](#)] [[INSPIRE](#)].
- [63] M.B. Gavela, R. Houtz, P. Quilez, R. Del Rey and O. Sumensari, *Flavor constraints on electroweak ALP couplings*, *Eur. Phys. J. C* **79** (2019) 369 [[arXiv:1901.02031](#)] [[INSPIRE](#)].
- [64] L. Merlo, F. Pobbe, S. Rigolin and O. Sumensari, *Revisiting the production of ALPs at B-factories*, *JHEP* **06** (2019) 091 [[arXiv:1905.03259](#)] [[INSPIRE](#)].
- [65] BELLE-II collaboration, *Search for Axion-Like Particles produced in e^+e^- collisions at Belle II*, *Phys. Rev. Lett.* **125** (2020) 161806 [[arXiv:2007.13071](#)] [[INSPIRE](#)].
- [66] D. Aristizabal Sierra, V. De Romeri, L.J. Flores and D.K. Papoulias, *Axionlike particles searches in reactor experiments*, *JHEP* **03** (2021) 294 [[arXiv:2010.15712](#)] [[INSPIRE](#)].
- [67] M.J. Dolan, F. Kahlhoefer, C. McCabe and K. Schmidt-Hoberg, *A taste of dark matter: Flavour constraints on pseudoscalar mediators*, *JHEP* **03** (2015) 171 [*Erratum ibid.* **07** (2015) 103] [[arXiv:1412.5174](#)] [[INSPIRE](#)].
- [68] B. Döbrich, F. Ertas, F. Kahlhoefer and T. Spadaro, *Model-independent bounds on light pseudoscalars from rare B-meson decays*, *Phys. Lett. B* **790** (2019) 537 [[arXiv:1810.11336](#)] [[INSPIRE](#)].
- [69] J. Beacham et al., *Physics Beyond Colliders at CERN: Beyond the Standard Model Working Group Report*, *J. Phys. G* **47** (2020) 010501 [[arXiv:1901.09966](#)] [[INSPIRE](#)].
- [70] S. Gori, G. Perez and K. Tobioka, *KOTO vs. NA62 Dark Scalar Searches*, *JHEP* **08** (2020) 110 [[arXiv:2005.05170](#)] [[INSPIRE](#)].

- [71] J. Martin Camalich, M. Pospelov, P.N.H. Vuong, R. Ziegler and J. Zupan, *Quark Flavor Phenomenology of the QCD Axion*, *Phys. Rev. D* **102** (2020) 015023 [[arXiv:2002.04623](#)] [[INSPIRE](#)].
- [72] CHARM collaboration, *Search for Axion Like Particle Production in 400 GeV Proton-Copper Interactions*, *Phys. Lett. B* **157** (1985) 458 [[INSPIRE](#)].
- [73] E.M. Riordan et al., *A Search for Short Lived Axions in an Electron Beam Dump Experiment*, *Phys. Rev. Lett.* **59** (1987) 755 [[INSPIRE](#)].
- [74] J.D. Bjorken et al., *Search for Neutral Metastable Penetrating Particles Produced in the SLAC Beam Dump*, *Phys. Rev. D* **38** (1988) 3375 [[INSPIRE](#)].
- [75] B. Döbrich, J. Jaeckel, F. Kahlhoefer, A. Ringwald and K. Schmidt-Hoberg, *ALPtraum: ALP production in proton beam dump experiments*, *JHEP* **02** (2016) 018 [[arXiv:1512.03069](#)] [[INSPIRE](#)].
- [76] S. Alekhin et al., *A facility to Search for Hidden Particles at the CERN SPS: the SHiP physics case*, *Rept. Prog. Phys.* **79** (2016) 124201 [[arXiv:1504.04855](#)] [[INSPIRE](#)].
- [77] B. Döbrich, J. Jaeckel and T. Spadaro, *Light in the beam dump. Axion-Like Particle production from decay photons in proton beam-dumps*, *JHEP* **05** (2019) 213 [*Erratum* *ibid.* **10** (2020) 046] [[arXiv:1904.02091](#)] [[INSPIRE](#)].
- [78] L. Darmé, F. Giacchino, E. Nardi and M. Raggi, *Invisible decays of axion-like particles: constraints and prospects*, *JHEP* **06** (2021) 009 [[arXiv:2012.07894](#)] [[INSPIRE](#)].
- [79] K.J. Kelly, S. Kumar and Z. Liu, *Heavy axion opportunities at the DUNE near detector*, *Phys. Rev. D* **103** (2021) 095002 [[arXiv:2011.05995](#)] [[INSPIRE](#)].
- [80] V. Brdar et al., *Axionlike Particles at Future Neutrino Experiments: Closing the Cosmological Triangle*, *Phys. Rev. Lett.* **126** (2021) 201801 [[arXiv:2011.07054](#)] [[INSPIRE](#)].
- [81] J.L. Feng, I. Galon, F. Kling and S. Trojanowski, *Axionlike particles at FASER: The LHC as a photon beam dump*, *Phys. Rev. D* **98** (2018) 055021 [[arXiv:1806.02348](#)] [[INSPIRE](#)].
- [82] G. Aielli et al., *Expression of interest for the CODEX-b detector*, *Eur. Phys. J. C* **80** (2020) 1177 [[arXiv:1911.00481](#)] [[INSPIRE](#)].
- [83] J.H. Chang, R. Essig and S.D. McDermott, *Supernova 1987A Constraints on Sub-GeV Dark Sectors, Millicharged Particles, the QCD Axion, and an Axion-like Particle*, *JHEP* **09** (2018) 051 [[arXiv:1803.00993](#)] [[INSPIRE](#)].
- [84] P. Carena, T. Fischer, M. Giannotti, G. Guo, G. Martínez-Pinedo and A. Mirizzi, *Improved axion emissivity from a supernova via nucleon-nucleon bremsstrahlung*, *JCAP* **10** (2019) 016 [*Erratum* *ibid.* **05** (2020) E01] [[arXiv:1906.11844](#)] [[INSPIRE](#)].
- [85] F. Ertas and F. Kahlhoefer, *On the interplay between astrophysical and laboratory probes of MeV-scale axion-like particles*, *JHEP* **07** (2020) 050 [[arXiv:2004.01193](#)] [[INSPIRE](#)].
- [86] G. Alonso-Álvarez, M.B. Gavela and P. Quilez, *Axion couplings to electroweak gauge bosons*, *Eur. Phys. J. C* **79** (2019) 223 [[arXiv:1811.05466](#)] [[INSPIRE](#)].
- [87] K. Choi, S.H. Im, C.B. Park and S. Yun, *Minimal Flavor Violation with Axion-like Particles*, *JHEP* **11** (2017) 070 [[arXiv:1708.00021](#)] [[INSPIRE](#)].
- [88] M. Chala, G. Guedes, M. Ramos and J. Santiago, *Running in the ALPs*, *Eur. Phys. J. C* **81** (2021) 181 [[arXiv:2012.09017](#)] [[INSPIRE](#)].

- [89] M. Bauer, M. Neubert, S. Renner, M. Schnubel and A. Thamm, *The Low-Energy Effective Theory of Axions and ALPs*, *JHEP* **04** (2021) 063 [[arXiv:2012.12272](#)] [[INSPIRE](#)].
- [90] R.D. Peccei and H.R. Quinn, *CP Conservation in the Presence of Instantons*, *Phys. Rev. Lett.* **38** (1977) 1440 [[INSPIRE](#)].
- [91] R.D. Peccei and H.R. Quinn, *Constraints Imposed by CP Conservation in the Presence of Instantons*, *Phys. Rev. D* **16** (1977) 1791 [[INSPIRE](#)].
- [92] S. Weinberg, *A New Light Boson?*, *Phys. Rev. Lett.* **40** (1978) 223 [[INSPIRE](#)].
- [93] F. Wilczek, *Problem of Strong P and T Invariance in the Presence of Instantons*, *Phys. Rev. Lett.* **40** (1978) 279 [[INSPIRE](#)].
- [94] J.E. Kim, *Weak Interaction Singlet and Strong CP Invariance*, *Phys. Rev. Lett.* **43** (1979) 103 [[INSPIRE](#)].
- [95] M.A. Shifman, A.I. Vainshtein and V.I. Zakharov, *Can Confinement Ensure Natural CP Invariance of Strong Interactions?*, *Nucl. Phys. B* **166** (1980) 493 [[INSPIRE](#)].
- [96] A.R. Zhitnitsky, *On Possible Suppression of the Axion Hadron Interactions* (in Russian), *Sov. J. Nucl. Phys.* **31** (1980) 260 [[INSPIRE](#)].
- [97] M. Dine, W. Fischler and M. Srednicki, *A Simple Solution to the Strong CP Problem with a Harmless Axion*, *Phys. Lett. B* **104** (1981) 199 [[INSPIRE](#)].
- [98] L. Calibbi, F. Goertz, D. Redigolo, R. Ziegler and J. Zupan, *Minimal axion model from flavor*, *Phys. Rev. D* **95** (2017) 095009 [[arXiv:1612.08040](#)] [[INSPIRE](#)].
- [99] Y. Ema, K. Hamaguchi, T. Moroi and K. Nakayama, *Flaxion: a minimal extension to solve puzzles in the standard model*, *JHEP* **01** (2017) 096 [[arXiv:1612.05492](#)] [[INSPIRE](#)].
- [100] C.D. Froggatt and H.B. Nielsen, *Hierarchy of Quark Masses, Cabibbo Angles and CP-violation*, *Nucl. Phys. B* **147** (1979) 277 [[INSPIRE](#)].
- [101] NA62 collaboration, *The Beam and detector of the NA62 experiment at CERN*, 2017 *JINST* **12** P05025 [[arXiv:1703.08501](#)] [[INSPIRE](#)].
- [102] KOTO collaboration, *Search for the $K_L \rightarrow \pi^0 \nu \bar{\nu}$ and $K_L \rightarrow \pi^0 X^0$ decays at the J-PARC KOTO experiment*, *Phys. Rev. Lett.* **122** (2019) 021802 [[arXiv:1810.09655](#)] [[INSPIRE](#)].
- [103] KLEVER PROJECT collaboration, *KLEVER: An experiment to measure $BR(K_L \rightarrow \pi^0 \nu \bar{\nu})$ at the CERN SPS*, [arXiv:1901.03099](#) [[INSPIRE](#)].
- [104] C. Vafa, *The String landscape and the swampland*, [hep-th/0509212](#) [[INSPIRE](#)].
- [105] B. Freivogel, T. Gasenzer, A. Hebecker and S. Leonhardt, *A Conjecture on the Minimal Size of Bound States*, *SciPost Phys.* **8** (2020) 058 [[arXiv:1912.09485](#)] [[INSPIRE](#)].
- [106] H. Georgi, D.B. Kaplan and L. Randall, *Manifesting the Invisible Axion at Low-energies*, *Phys. Lett. B* **169** (1986) 73 [[INSPIRE](#)].
- [107] G. Grilli di Cortona, E. Hardy, J. Pardo Vega and G. Villadoro, *The QCD axion, precisely*, *JHEP* **01** (2016) 034 [[arXiv:1511.02867](#)] [[INSPIRE](#)].
- [108] P. Di Vecchia and G. Veneziano, *Chiral Dynamics in the Large N Limit*, *Nucl. Phys. B* **171** (1980) 253 [[INSPIRE](#)].
- [109] C. Cornella, P. Paradisi and O. Sumensari, *Hunting for ALPs with Lepton Flavor Violation*, *JHEP* **01** (2020) 158 [[arXiv:1911.06279](#)] [[INSPIRE](#)].

- [110] P. Escribano and A. Vicente, *Ultralight scalars in leptonic observables*, *JHEP* **03** (2021) 240 [[arXiv:2008.01099](#)] [[INSPIRE](#)].
- [111] J. Ellis, *TikZ-Feynman: Feynman diagrams with TikZ*, *Comput. Phys. Commun.* **210** (2017) 103 [[arXiv:1601.05437](#)] [[INSPIRE](#)].
- [112] N. Carrasco, P. Lami, V. Lubicz, L. Riggio, S. Simula and C. Tarantino, *$K \rightarrow \pi$ semileptonic form factors with $N_f = 2 + 1 + 1$ twisted mass fermions*, *Phys. Rev. D* **93** (2016) 114512 [[arXiv:1602.04113](#)] [[INSPIRE](#)].
- [113] FLAVOUR LATTICE AVERAGING GROUP collaboration, *FLAG Review 2019: Flavour Lattice Averaging Group (FLAG)*, *Eur. Phys. J. C* **80** (2020) 113 [[arXiv:1902.08191](#)] [[INSPIRE](#)].
- [114] B. Batell, M. Pospelov and A. Ritz, *Multi-lepton Signatures of a Hidden Sector in Rare B Decays*, *Phys. Rev. D* **83** (2011) 054005 [[arXiv:0911.4938](#)] [[INSPIRE](#)].
- [115] W.A. Bardeen, R.D. Peccei and T. Yanagida, *Constraints on variant axion models*, *Nucl. Phys. B* **279** (1987) 401 [[INSPIRE](#)].
- [116] E949 and E787 collaborations, *Measurement of the $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ branching ratio*, *Phys. Rev. D* **77** (2008) 052003 [[arXiv:0709.1000](#)] [[INSPIRE](#)].
- [117] NA62 collaboration, *The NA62 experiment at CERN: status and perspectives*, in *12th Conference on Flavor Physics and CP-violation*, (2014) [[arXiv:1407.8213](#)] [[INSPIRE](#)].
- [118] NA62 collaboration, *Search for a feebly interacting particle X in the decay $K^+ \rightarrow \pi^+ X$* , *JHEP* **03** (2021) 058 [[arXiv:2011.11329](#)] [[INSPIRE](#)].
- [119] S. Oda, Y. Shoji and D.-S. Takahashi, *High Scale Validity of the DFSZ Axion Model with Precision*, *JHEP* **03** (2020) 011 [[arXiv:1912.01147](#)] [[INSPIRE](#)].
- [120] F. Kling, S. Su and W. Su, *2HDM Neutral Scalars under the LHC*, *JHEP* **06** (2020) 163 [[arXiv:2004.04172](#)] [[INSPIRE](#)].
- [121] L. Di Luzio, M. Giannotti, E. Nardi and L. Visinelli, *The landscape of QCD axion models*, *Phys. Rept.* **870** (2020) 1 [[arXiv:2003.01100](#)] [[INSPIRE](#)].
- [122] J. Quevillon and C. Smith, *Axions are blind to anomalies*, *Eur. Phys. J. C* **79** (2019) 822 [[arXiv:1903.12559](#)] [[INSPIRE](#)].
- [123] S.L. Adler, *Axial vector vertex in spinor electrodynamics*, *Phys. Rev.* **177** (1969) 2426 [[INSPIRE](#)].
- [124] J.S. Bell and R. Jackiw, *A PCAC puzzle: $\pi^0 \rightarrow \gamma\gamma$ in the σ model*, *Nuovo Cim. A* **60** (1969) 47 [[INSPIRE](#)].
- [125] W.A. Bardeen, *Anomalous Ward identities in spinor field theories*, *Phys. Rev.* **184** (1969) 1848 [[INSPIRE](#)].
- [126] H.H. Patel, *Package-X: A Mathematica package for the analytic calculation of one-loop integrals*, *Comput. Phys. Commun.* **197** (2015) 276 [[arXiv:1503.01469](#)] [[INSPIRE](#)].
- [127] L.J. Hall and M.B. Wise, *Flavor changing Higgs-boson couplings*, *Nucl. Phys. B* **187** (1981) 397 [[INSPIRE](#)].
- [128] C.A.J. O'Hare and E. Vitagliano, *Cornering the axion with CP-violating interactions*, *Phys. Rev. D* **102** (2020) 115026 [[arXiv:2010.03889](#)] [[INSPIRE](#)].
- [129] L. Di Luzio, R. Gröber and P. Paradisi, *Hunting for the CP-violating ALP*, [[arXiv:2010.13760](#)] [[INSPIRE](#)].

- [130] S. Dar, *The Neutron EDM in the SM: A Review*, [hep-ph/0008248](#) [INSPIRE].
- [131] nEDM collaboration, *Measurement of the permanent electric dipole moment of the neutron*, *Phys. Rev. Lett.* **124** (2020) 081803 [[arXiv:2001.11966](#)] [INSPIRE].
- [132] L. Di Luzio, F. Mescia and E. Nardi, *Window for preferred axion models*, *Phys. Rev. D* **96** (2017) 075003 [[arXiv:1705.05370](#)] [INSPIRE].
- [133] A. Ayala, I. Domínguez, M. Giannotti, A. Mirizzi and O. Straniero, *Revisiting the bound on axion-photon coupling from Globular Clusters*, *Phys. Rev. Lett.* **113** (2014) 191302 [[arXiv:1406.6053](#)] [INSPIRE].
- [134] F. Capozzi and G. Raffelt, *Axion and neutrino bounds improved with new calibrations of the tip of the red-giant branch using geometric distance determinations*, *Phys. Rev. D* **102** (2020) 083007 [[arXiv:2007.03694](#)] [INSPIRE].
- [135] M. Giannotti, I. Irastorza, J. Redondo and A. Ringwald, *Cool WISPs for stellar cooling excesses*, *JCAP* **05** (2016) 057 [[arXiv:1512.08108](#)] [INSPIRE].
- [136] S. Hoof, F. Kahlhoefer, P. Scott, C. Weniger and M. White, *Axion global fits with Peccei-Quinn symmetry breaking before inflation using GAMBIT*, *JHEP* **03** (2019) 191 [Erratum *ibid.* **11** (2019) 099] [[arXiv:1810.07192](#)] [INSPIRE].
- [137] M.M. Miller Bertolami, B.E. Melendez, L.G. Althaus and J. Isern, *Revisiting the axion bounds from the Galactic white dwarf luminosity function*, *JCAP* **10** (2014) 069 [[arXiv:1406.7712](#)] [INSPIRE].
- [138] M. Giannotti, I.G. Irastorza, J. Redondo, A. Ringwald and K. Saikawa, *Stellar Recipes for Axion Hunters*, *JCAP* **10** (2017) 010 [[arXiv:1708.02111](#)] [INSPIRE].
- [139] M. Srednicki, *Axion Couplings to Matter. 1. CP Conserving Parts*, *Nucl. Phys. B* **260** (1985) 689 [INSPIRE].
- [140] T. Vonk, F.-K. Guo and U.-G. Meißner, *Precision calculation of the axion-nucleon coupling in chiral perturbation theory*, *JHEP* **03** (2020) 138 [[arXiv:2001.05327](#)] [INSPIRE].
- [141] N. Bar, K. Blum and G. D’Amico, *Is there a supernova bound on axions?*, *Phys. Rev. D* **101** (2020) 123025 [[arXiv:1907.05020](#)] [INSPIRE].
- [142] P. Cigan et al., *High angular resolution ALMA images of dust and molecules in the SN 1987A ejecta*, *Astrophys. J.* **886** (2019) 51 [[arXiv:1910.02960](#)] [INSPIRE].
- [143] D. Page, M.V. Beznogov, I. Garibay, J.M. Lattimer, M. Prakash and H.-T. Janka, *NS 1987A in SN 1987A*, *Astrophys. J.* **898** (2020) 125 [[arXiv:2004.06078](#)] [INSPIRE].
- [144] A. Ringwald and K. Saikawa, *Axion dark matter in the post-inflationary Peccei-Quinn symmetry breaking scenario*, *Phys. Rev. D* **93** (2016) 085031 [Addendum *ibid.* **94** (2016) 049908] [[arXiv:1512.06436](#)] [INSPIRE].
- [145] C. Bobeth, T. Ewerth, F. Krüger and J. Urban, *Analysis of neutral Higgs boson contributions to the decays $\bar{B}_s \rightarrow \ell^+ \ell^-$ and $\bar{B} \rightarrow K \ell^+ \ell^-$* , *Phys. Rev. D* **64** (2001) 074014 [[hep-ph/0104284](#)] [INSPIRE].
- [146] H.E. Logan and U. Nierste, *$B_{s,d} \rightarrow \ell^+ \ell^-$ in a two Higgs doublet model*, *Nucl. Phys. B* **586** (2000) 39 [[hep-ph/0004139](#)] [INSPIRE].