

Vectorial axion-down-strange coupling from $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ data

Diego Guadagnoli^{1,*}, Axel Iohner^{1,†}, Cristina Lazzeroni^{2,‡}, Diego Martínez Santos^{3,§},
Joel C. Swallow^{4,||} and Claudio Toni^{1,¶}

¹*LAPTh, Université Savoie Mont-Blanc et CNRS, 74941 Annecy, France*

²*School of Physics and Astronomy, University of Birmingham, Birmingham, United Kingdom*

³*Universidad de A Coruña, Ferrol Industrial Campus, Dr. Vázquez Cabrera, s/n 15403, A Coruña, Spain*

⁴*CERN, Esplanade des Particules 1, 1211 Meyrin, Switzerland*



(Received 31 March 2025; revised 6 March 2026; accepted 6 July 2026; published 2 September 2026)

We reinterpret publicly available $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ data collected by NA62 from 2016 to 2024 to constrain the fundamental vectorial coupling of the QCD axion to down and strange quarks. Using a fully reproducible likelihood analysis and a complete renormalization-group evolution of the axion couplings from the Peccei-Quinn (PQ) scale to the kaon scale, we translate the experimental limit into bounds on both the low-energy flavor-violating coupling and the fundamental UV parameters. In the generic regime where strong contributions dominate the decay amplitude, we obtain $|(F_V)_{sd}(\mu_K)| > 1.6 \times 10^{12} \text{ GeV}$. The coexistence of parametrically suppressed weak contributions implies a second, conceptually distinct constraint: The fact that weak-amplitude dominance arises from highly tuned UV coupling configurations translates into a conservative general lower limit on the PQ scale, $f_a > 4.9 \times 10^4 \text{ GeV}$. These results provide the strongest accelerator-based constraints on axion-induced $d \leftrightarrow s$ transitions and establish a robust lower bound on f_a , complementary to astrophysical limits.

DOI: [10.1103/r2rg-mc5s](https://doi.org/10.1103/r2rg-mc5s)

I. INTRODUCTION

New spin-zero, odd-parity particles below the GeV scale—collectively referred to as axions or axionlike particles (ALPs)—are hypothetical states, searched for by an expanding range of laboratory experiments [1–3], astrophysical observations [4–6], cosmological probes [7], and phenomenological considerations (see, e.g., Refs. [8–10]).

A well-defined thread within this broader context is the search for the so-called “invisible” [11–14] QCD axion [15–18], hereafter simply the axion. This is a highly plausible particle, as it provides a uniquely elegant resolution to the strong- CP problem and a compelling framework for accounting for the fraction of the universe’s total density attributed to dark matter [19–21]. Unlike generic ALPs, the axion’s mass m_a and decay constant f_a are intrinsically related: Smaller values of m_a correspond to weaker interactions with standard matter, with couplings suppressed parametrically as $1/f_a$.

Current constraints imply that $m_a \ll m_\pi$ [9]. Despite the accordingly feeble couplings, the small mass means that axions may be produced, potentially even copiously, within controlled accelerator environments, without the need to push the energy frontier.

Among the many experimental approaches to probing the existence of axions, one of the most direct is their potential production in decays recorded at high-intensity facilities such as NA62. This approach is well motivated for several reasons. First, this method has for decades been used to study decays to—and even to unambiguously identify [22]—neutrinos, now-established feebly interacting, uncharged, near-massless particles with properties largely analogous to those of axions. Second, from a theoretical point of view, light meson decays involving an axion admit a systematic, effective-theory description, known as chiral perturbation theory (ChPT) [23,24], which ensures reliable predictions with a controlled uncertainty. Axion interactions with light mesons can be included in this description by a careful generalization [25] of the same logic that leads to the construction of ChPT. In this construction, starting from the Lagrangian terms

$$\mathcal{L}_a \supset \frac{\partial_\mu a}{2f_a} (\bar{q} \gamma^\mu \mathbf{k}_V q + \bar{q} \gamma^\mu \gamma_5 \mathbf{k}_A q) - \frac{\alpha_s}{8\pi} \frac{a}{f_a} G\tilde{G}, \quad (1)$$

with $q = (u, d, s)^T$ [26], the ChPT couplings to hadrons are unambiguously determined once the $\mathbf{k}_{V,A}$ coupling matrices are fixed—up to the standard hadronic low-energy constants at the chosen chiral order and to an unphysical reparametrization ambiguity [25,27].

*Contact author: diego.guadagnoli@lapth.cnrs.fr

†Contact author: axel.iohner@lapth.cnrs.fr

‡Contact author: cristina.lazzeroni@cern.ch

§Contact author: diego.martinez.santos@cern.ch

||Contact author: joel.swallow@cern.ch

¶Contact author: claudio.toni@lapth.cnrs.fr

The axion-down-strange couplings $(k_{V,A})_{sd}$ —the only allowed off-diagonal entries in Eq. (1)—are thus two fundamental couplings of the QCD axion to matter. Given that they generate $d \leftrightarrow s$ transitions otherwise highly suppressed in the standard model (SM), these couplings can lead to “very dramatic” effects [25]. The strongest probe of the vectorial $|(k_V)_{sd}|$ are $K \rightarrow \pi a$ decays. The diagonal axial couplings $(k_A)_{uu,dd}$ (reals) are best probed by neutron stars [28]. Finally, it has recently been shown that supernova cooling induced by axions radiated from strange matter in the star core [29,30] is sensitive to the axial $|(k_A)_{sd}|$ and $(k_A)_{ss}$.

Here, we probe $K^+ \rightarrow \pi^+ a$ by reinterpreting experimental studies of the ultrarare decay $K^+ \rightarrow \pi^+ \nu \bar{\nu}$, which provides one of the most exquisite precision tests of the SM. The tiny $O(10^{-10})$ branching fraction of the K^+ into this decay mode [31–33] is dominated by short-distance contributions, and the necessary hadronic matrix element can be extracted from data using isospin. The decay is thus uniquely rare and uniquely clean theoretically. The NA62 experiment at CERN was designed with the measurement of $\mathcal{B}(K^+ \rightarrow \pi^+ \nu \bar{\nu}) \equiv \mathcal{B}_{\nu\bar{\nu}}$ as its primary goal. Using the decay-in-flight technique, it has reported a sequence of increasingly precise measurements based on the 2016–2018 data [34–36], the 2021–2022 dataset [37], the full 2016–2022 sample [38], and most recently the 2023–2024 data [39].

Using all these measurements, one can constrain decays to π^+ plus invisible states, in particular, $K^+ \rightarrow \pi^+ a$. Such a signal would appear as an excess over $K^+ \rightarrow \pi^+ \nu \bar{\nu}$. At an intuitive level, the experimental uncertainty $\delta\mathcal{B}_{\nu\bar{\nu}}$ quantifies the available room for this excess and thus sets the expected sensitivity to $\mathcal{B}(K^+ \rightarrow \pi^+ a)$. The combination of 2016–2024 NA62 data yields $\delta\mathcal{B}_{\nu\bar{\nu}} \simeq 1.9 \times 10^{-11}$, predominantly statistical. Our first result is a fully reproducible likelihood-based reinterpretation relying exclusively on public data, which quantitatively confirms this expectation.

We apply this method to infer the strongest existing bound on $|(k_V)_{sd}|/f_a$ at the kaon scale—our second result, obtained from the entire 2016–2024 dataset analyzed so far. This result will remain state of the art until the 2025–2026 data, currently being collected by NA62, are analyzed. Using a full renormalization-group evolution from the Peccei-Quinn scale to hadronic scales, we express this constraint directly in terms of the fundamental UV couplings and the PQ scale f_a . This allows us to expose unambiguously the hierarchy between the strong and parametrically suppressed weak contributions to the amplitude. Accordingly, two conceptually distinct bounds emerge. In the generic regime of strong-amplitude dominance, we obtain the leading constraint. Conversely, weak-amplitude dominance requires highly tuned UV coupling configurations. This tuning requirement allows the weak contribution to define a conservative general floor on the PQ scale, extracted from controlled accelerator data. This forms our third result.

II. REINTERPRETATION TECHNIQUE

To bound the $K^+ \rightarrow \pi^+ a$ decay through the $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ measurements at NA62, we build on the procedure of Ref. [40]. We perform a fully frequentist hypothesis test using a shape analysis of the invisible invariant mass and an unbinned profile likelihood ratio test statistic, based solely

on public data. The likelihood function is constructed as the product of three contributions, i.e., $\mathcal{L} = \mathcal{L}_1 \mathcal{L}_2 \mathcal{L}_3$.

The first term is a Poisson distribution for the total number of events n_{tot} with expectation value fixed to the total number of observed events n_{obs} . The variable n_{tot} can be decomposed as $n_{\text{tot}} = n_b + n_a$, with $n_{b,a}$ being the number of background and axion events, respectively.

The number n_a is related to the branching ratio (BR) as $\mathcal{B}(K^+ \rightarrow \pi^+ a) = n_a \times \mathcal{B}_{\text{SES}}$, where $\mathcal{B}_{\text{SES}}$ denotes the associated single-event sensitivity (SES). $\mathcal{B}_{\text{SES}} = \mathcal{B}_{\text{SES}}(m_a)$ is provided in, e.g., the rightmost panel of Fig. 2 of Ref. [40], which pertains to the 2017 data. Limits for other datasets are obtained by rescaling with the ratio of the corresponding $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ SES values. Note that n_b also includes the $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ events, assuming that $\mathcal{B}_{\nu\bar{\nu}}^{\text{SM}} = 8.4 \times 10^{-11}$ as in Refs. [37, 41].

Given n_{tot} , its two components are distributed as a function of the invisible invariant mass squared m_{inv}^2 according to a multinomial likelihood, yielding $\mathcal{L}_2$ as

$$\mathcal{L}_2 = \prod_{j=1}^{n_{\text{obs}}} \left[\frac{n_b}{n_{\text{tot}}} g_b(m_j^2) + \frac{n_a}{n_{\text{tot}}} g_a(m_j^2) \right], \quad (2)$$

where $g_b(m_{\text{inv}}^2)$ is a function reproducing the distribution of the background events and normalized to unit integral over the signal regions. We took this function from published data, i.e., from the right panel of Fig. 10 in Ref. [37]. Conversely, $g_a(m_{\text{inv}}^2)$ is a normal distribution centered at the axion mass value with uncertainty given by the invariant mass resolution at fixed m_a , $\sigma_{m_a^2}$, taken from the left panel of Fig. 2 of Ref. [40]. Finally, the last likelihood factor, $\mathcal{L}_3$, is a Poisson-distributed constraint term for the background yield, written as [42]

$$\mathcal{L}_3 = \frac{(\tau n_b)^{n_{\text{off}}}}{n_{\text{off}}!} e^{-\tau n_b}, \quad (3)$$

with $\tau = \mu_b/\sigma_b^2$ and $n_{\text{off}} = (\mu_b/\sigma_b)^2$ constructed from the estimated mean μ_b and uncertainty σ_b . We neglect the uncertainty on the rest of the nuisance parameters considered in Ref. [40], namely, the SES uncertainty, as well as the error on the parameters of the axion mass probability density function. By inspection, these uncertainties are subleading with respect to those discussed, and are expected to affect our results in a minor way. Minimizing $\chi^2(n_b, n_a) \equiv -2 \ln \mathcal{L}$, and demanding $\Delta\chi^2 \equiv \chi^2 - \chi_{\text{min}}^2 = \chi^2(90\%, 2 \text{ d.o.f.})$, we obtain a limit on $\mathcal{B}(K^+ \rightarrow \pi^+ a)$ as a function of m_a . We tested this procedure by comparing our obtained limit with the NA62 results based on 2016–2018 [36] and 2017 data [40], respectively, finding very good agreement. The details of this validation, and plots quantifying the agreement, are provided in Appendix A. We apply this procedure to two datasets: the full published NA62 dataset, encompassing data from 2016 to 2022, and the dataset further including the data recorded in 2023–2024 [39]. We obtain the two BR limits shown in Fig. 1 (left). At $m_a = 0$, these limits read

$$\mathcal{B}(K^+ \rightarrow \pi^+ a) < 2.8 \times 10^{-11} \equiv \mathcal{B}_{16-22}, \quad (4a)$$

$$\mathcal{B}(K^+ \rightarrow \pi^+ a) < 1.4 \times 10^{-11} \equiv \mathcal{B}_{16-24}, \quad (4b)$$

at 90% C.L. for the two mentioned datasets. The bounds [Eq. (4)] improve over Refs. [36,38] by a factor of about 2




FIG. 1. (Left) Upper limit on $\mathcal{B}(K^+ \rightarrow \pi^+ a)$ from NA62 using the published 2016–2022 dataset [38], and the extended 2016–2024 sample including the new 2023–2024 data [39]. The gray dashed line shows $\delta\mathcal{B}_{v\bar{v}}$. The range [100, 161] MeV is excluded due to $K^+ \rightarrow \pi^+ \pi^0$ background contamination [37,40]. (Right) Corresponding lower bound on the flavor-violating-coupling rescaled Peccei-Quinn scale (see text for details).

in both cases, in line with the corresponding increase in the data sample size.

Reference [38], which appeared months after the first version of this study, includes an analysis of $K^+ \rightarrow \pi^+ X$, where X is an ALP. At $m_X = 0$, Ref. [38] agrees with our result in Eq. (4a) to within one unit of the last significant digit. This agreement constitutes a strong validation of our approach: It confirms the possibility of retrieving the leading uncertainties from public data alone, our treatment of these uncertainties, and the negligible impact of those we did not include. While a full NA62 search remains the benchmark, the above comparison shows that our approach provides an excellent approximation and has the further key advantage of being fully reproducible. Below, we use the state-of-the-art limit [Eq. (4b)] to set unprecedented bounds on the relevant axion coupling or, alternatively, on the axion scale.

III. THEORY PREDICTION

The axion, as the near-Goldstone boson of a global $U(1)$ Peccei-Quinn symmetry, can be incorporated into ChPT as elucidated in Ref. [25]. We refer to the $O(p^2)$ ChPT Lagrangian augmented with an axion as the “strong” Lagrangian $\mathcal{L}_s$. A generalization of this procedure enables inclusion of the effect of strangeness (S) changing weak interactions. We denote as $\mathcal{L}_w$ the $|\Delta S| = 1$ contributions that are enhanced by the so-called $\Delta I = 1/2$ rule. These include the well-known G_8 term and, in the presence of an axion, the so-called G_8^θ term [43]. We find that the coupling G_8^θ is not independent of G_8 . Requiring that the $K^0 \rightarrow Z$ amplitude vanishes to the perturbative order $p^2 G_F$ we are considering [44] implies the relation $G_8^\theta = -2G_8$ up to loop-suppressed penguin contributions. We are thus able to include the G_8^θ contribution in a fully quantitative way for the first time. The amplitude $\mathcal{M}$ for the $K^+ \rightarrow \pi^+ a$ decay can be decomposed into two pieces,

$$\mathcal{M}(\mathbf{k}) \equiv \mathcal{M}_s((k_V)_{sd}) + \mathcal{M}_w((k_A)_{uu,dd,ss}, (k_V)_{dd,ss}), \quad (5)$$

induced by $\mathcal{L}_s$ and $\mathcal{L}_w$, respectively. We collect explicit formulas for the Lagrangians and for the amplitudes in Appendix B.

The k couplings in Eq. (5) are understood at a renormalization scale $\mu_K \approx m_K$. A bound from $K \rightarrow \pi a$ on these couplings is informative on the strength of the axion-induced flavor-violating mechanism at the energy of the decay, but it says little about the fundamental couplings at the high scale $\mu_{PQ} \sim 4\pi f_a$ where the PQ symmetry is broken. To this end, one needs to renormalization-group (RG) evolve the k couplings appearing in Eq. (5) to μ_{PQ} .

We note that axion-induced flavor violation at low energies can arise not only from the same coupling at high energies, but also from the misalignment between the up- and down-type quark Yukawa couplings $Y_{u,d}$, parametrized by the Cabibbo-Kobayashi-Maskawa matrix V . Both these inputs, therefore, have to be carefully included. Using $Y_{u,d}$ and V at the electroweak scale $\mu_{EW} \sim m_t$ (e.g., Ref. [45]) as boundary conditions, we obtain their counterparts at the μ_{PQ} scale using Ref. [46]. We then RG evolve from μ_{PQ} to μ_{EW} the coupled equations (implemented following the comprehensive Ref. [47]) for Yukawa and axion couplings, using the PQ-scale Yukawa and axion couplings as boundary conditions. The axion couplings amount to 35 degrees of freedom (d.o.f.) in total and include the following: the axion-gauge boson couplings c_{BB} and c_{WW} [48]; the axion couplings to the chiral matter multiplets $\Phi = Q, L$ and u, d, e , denoted as $(k_\Phi)_{ij}$, with i, j generation indices—that we distinguish from the light-quark indices used, e.g., around Eq. (1), by denoting the latter indices with light-quark symbols.

Importantly, the definition of k -coupling entries at the PQ scale is quark-field-basis dependent; i.e., flavor-violating effects can be shuffled from nondiagonal Yukawa (i.e., quark-mass) matrices to nondiagonal k -coupling matrices. Given that our underlying transition is $s \rightarrow d$, the most convenient basis in our case is the down-quark basis, where the down-type Yukawa matrix is diagonal. In this basis, PQ-scale dynamics is thus understood to be flavor-conserving or violating in the down sector [49] and $Y_d(\mu_{PQ})$ is, by definition, diagonal.

To the thus-obtained Yukawa and axion couplings at the μ_{EW} scale, we add matching contributions at μ_{EW} [47].

Finally, RG effects below μ_{EW} are, depending on the coupling, either absent or tiny, and can be safely neglected [50].

IV. RESULTS

We next discuss the implications of the relation $\mathcal{B}(|\mathcal{M}|) \leq \mathcal{B}_{\text{exp}}$, with $\mathcal{M}$ from Eq. (5). Here, $\mathcal{M}_s$ is proportional to the relevant flavor-violating axion coupling $(k_V)_{\text{sd}}(\mu_K)$. Conversely, $\mathcal{M}_w$ depends on flavor-conserving k couplings only and has an overall parametric suppression of order $G_F F_0^2 \approx 1.0 \times 10^{-7}$. Therefore, the natural expectation is that strong contributions to $K \rightarrow \pi a$ are larger than weak contributions, i.e., $|\mathcal{M}_s| \gg |\mathcal{M}_w|$, in which case the experimental bounds [Eq. (4)] translate into bounds on $|(k_V)_{\text{sd}}(\mu_K)|/(2f_a)$ or, equivalently, its inverse, that for $m_a = 0$ read

$$|(F_V)_{\text{sd}}(\mu_K)| > 1.1 \times 10^{12} \text{ GeV}, \quad (6a)$$

$$|(F_V)_{\text{sd}}(\mu_K)| > 1.6 \times 10^{12} \text{ GeV}. \quad (6b)$$

These two bounds are 1.6 and 2.4 times, respectively, stronger than the latest data-driven bound [51] on the effective scale of the underlying interaction. In Fig. 1 (right), we show a generalization of Eq. (6) for any m_a in the range of values probed by NA62.

Our Fig. 1 assumes a π^+ plus missing-energy final state, i.e., that the axion escapes the approximately 130 m-long instrumented region downstream of the kaon decay without decaying visibly. In the mass range shown in the figure, the axion could in principle decay into $\gamma\gamma$, e^+e^- , or $\mu^+\mu^-$ (decays into hadrons require the two-pion threshold). For $m_a \lesssim 1$ MeV, only the $a \rightarrow \gamma\gamma$ channel is kinematically allowed. A sufficient condition for the axion to escape detection within the instrumented region is $|g_{a\gamma\gamma}| \lesssim 10^{-8}$, which ensures a decay length larger than the detector size. Existing bounds on $g_{a\gamma\gamma}$ [52] are significantly stronger than this value, so our assumption is safely satisfied in the regime $m_a \lesssim 1$ MeV. For $m_a \gtrsim 1$ MeV, the m_a - f_a relation is strongly model dependent. Rather than entering model-specific considerations, we provide the data-driven requirement $\tau_a \gtrsim 10$ ns [38], which can be imposed on any given model.

In QCD-axion models, m_a is uniquely determined by f_a . Expressing f_a through $(F_V)_{\text{sd}}$ and $(k_V)_{\text{sd}}$, and inserting the bound [Eq. (6)], one finds that for a natural coupling $|(k_V)_{\text{sd}}| \sim 1$ the corresponding axion mass lies in the range $m_a \sim 10^{-15}$ to 10^{-14} GeV. Reducing $|(k_V)_{\text{sd}}|$ to $O(10^{-7})$ [i.e., $O(G_F F_0^2)$] raises the implied mass to $m_a \sim 10^{-7}$ GeV (see also Appendix C). For even smaller values of $|(k_V)_{\text{sd}}|$, the strong and weak contributions to the amplitude become comparable and one enters the finely tuned regime discussed below. It follows that over the relevant parameter range the axion is effectively massless on accelerator scales: Both the dynamical scale of the process ($m_K \sim 500$ MeV) and the experimental resolution exceed the implied QCD-axion mass by many orders of magnitude. Although accelerator limits are therefore insensitive to such small m_a values, they impose strong constraints on the flavor-violating coupling and thereby indirectly single out a well-defined QCD-axion mass window.

Based on the discussion after Eq. (5), we can also study the inequality $\mathcal{B}(|\mathcal{M}|) \leq \mathcal{B}_{\text{exp}}$ directly in terms of the UV couplings at the scale μ_{PQ} . Since the amplitude is linear in

$1/f_a$, this inequality can be rewritten as

$$f_a \geq \sqrt{\frac{\mathcal{B}(f_a \cdot |\mathcal{M}|)}{\mathcal{B}_{\text{exp}}}}, \quad (7)$$

where the right-hand side depends on the UV-scale k couplings only—i.e., it does not depend on f_a . This leads to

$$f_a \geq \underbrace{4.85 \times 10^4 \text{ GeV}}_{\text{defines } \tilde{f}_a} \cdot \underbrace{|1 + \alpha \cdot \mathbf{c}|}_{\text{defines } |z|}, \quad (8)$$

where $\mathbf{c}$ is the vector of 35 couplings at μ_{PQ} and α is the vector of multiplying coefficients. The unit term inside the absolute-value factor in Eq. (8) tracks the contribution to $\mathcal{B}_a$ due to the $aG\tilde{G}$ term in Eq. (1): The $c_i \rightarrow 0$ limit yields the bound $f_a \geq \tilde{f}_a$.

The factor $|z|$ in Eq. (8) reflects the coexistence of strong and weak contributions, the latter being parametrically suppressed by $O(10^{-7})$. Accordingly, the coefficients α_i span several orders of magnitude, from $O(10^{-2})$ up to $O(10^6)$ (a complete tabulation at the reference scale $\mu_{\text{PQ}} = 10^{12}$ GeV is provided in Table I). No established theory fixes the UV values of the couplings c_i , but they plausibly lie in the range $[-5, 5]$. (Larger $|c_i|$ would raise perturbativity concerns; moreover, the largest SM Yukawa coupling—conceptually similar to the couplings in Eq. (1)—is ~ 1 .) For generic c_i values in this range, one finds that $|z| \sim |z|_{\text{peak}} \equiv 1 \cdot 10^8$, implying typical bounds $f_a^{\text{gen}} \sim 10^{12}$ GeV. This becomes $f_a^{\text{FC}} \sim 10^8$ GeV if one assumes flavor-violating couplings to be zero at the PQ scale, representing theories where the axion is flavor-conserving at that scale. In both cases, illustrated in Fig. 4, the f_a bound is thus orders of magnitude stronger than $\tilde{f}_a$.

Achieving $|z| \sim O(1)$ (weak-amplitude dominance) requires cancellations among many complex contributions: Geometrically, it corresponds to forcing a complex quantity whose typical magnitude is $\sim 10^7$ to lie inside a disk of radius $O(1)$ around the origin. For random UV choices $c_i \in [-5, 5]$, this region has an exceedingly small measure. For $r \ll |z|_{\text{peak}}$, the small- $|z|$ probability scales as r^2 , which allows an extrapolation of the Monte Carlo results described in Appendix D 2. We obtain $P(|z| \leq 1) \sim 10^{-16}$, implying that $O(10^{16})$ trials are required for a single instance of $|z| \lesssim 1$.

We therefore infer the bound

$$f_a \geq \tilde{f}_a = 4.9 \times 10^4 \text{ GeV}, \quad (9)$$

which should be regarded as a general and conservative lower limit: In the vast majority of parameter space, the implied bound on f_a is much stronger. The limit in Eq. (9) corresponds to $m_a \leq 1.3 \times 10^{-7}$ GeV.

Applying to the weak-amplitude-dominated case, the bound in Eq. (9) is sensitive to the choice of weak couplings g_8 and G_8^θ —See Appendix B for their definition and the reference values used in this work. For example, in the absence of G_8^θ [43], whose effect we are in a position to quantitatively include for the first time in this work, and using either the “classic” $g_8 \simeq 5$ or instead 3.61 [53], one would obtain 1.70×10^5 or 1.35×10^5 GeV, respectively.

While $f_a \geq \tilde{f}_a$ is a conservative lower limit, we also examined explicitly two specific corners of parameter space in which the bound in Eq. (9) can be approached: single-coupling tuning, and engineered cancellations between one

TABLE I. Coefficients α_i of the linear form $z = 1 + \alpha \cdot \mathbf{c}$ at the reference scale $\mu_{\text{PQ}} = 10^{12}$ GeV, ordered by decreasing magnitude $|\alpha_i|$. Phases are given in radians.

c_i	$ \alpha_i $	$\arg(\alpha_i)$
c_{d12}^{im}	4.65×10^6	+1.678
c_{d12}^{re}	4.65×10^6	-1.463
c_{Q12}^{im}	4.65×10^6	+1.678
c_{Q12}^{re}	4.65×10^6	-1.463
c_{Q13}^{im}	2.78×10^4	+1.482
c_{Q13}^{re}	2.78×10^4	-1.659
c_{Q23}^{re}	5.86×10^3	+0.281
c_{Q23}^{im}	5.86×10^3	-1.292
c_{d13}^{re}	8.27×10^2	-0.089
c_{d13}^{im}	8.27×10^2	+1.482
c_{Q22}	2.70×10^2	-2.848
c_{Q11}	1.16×10^2	+0.293
c_{u33}	8.56×10^1	+0.300
c_{Q33}	7.00×10^1	+0.300
c_{d22}	2.02×10^1	-2.842
c_{d11}	1.90×10^1	+0.311
c_{d23}^{re}	9.06	+0.281
c_{d23}^{im}	9.06	-1.292
c_{u23}^{re}	1.98×10^1	-3.055
c_{u23}^{im}	3.81	-2.807
c_{u22}	0.733	-0.260
c_{u11}	0.995	-3.137
c_{d33}	0.269	-2.849
c_{WW}	0.104	-2.843
c_{u13}^{im}	0.0893	+1.499
c_{u13}^{re}	0.0815	-0.107
c_{L11}	0.0331	+0.296
c_{L22}	0.0331	+0.296
c_{L33}	0.0311	+0.296
c_{BB}	7.73×10^{-3}	-2.845
c_{e11}	7.73×10^{-3}	-2.845
c_{e22}	7.72×10^{-3}	-2.846
c_{e33}	5.79×10^{-3}	-2.844
c_{u12}^{im}	8.42×10^{-3}	+1.678
c_{u12}^{re}	7.57×10^{-3}	-3.035

parameter and the sum of the rest. A detailed analysis is presented in Appendix D3. We find that these configurations occupy a negligibly small region of the allowed UV parameter space and, moreover, require relations connecting quantities that are structurally unrelated, such as UV couplings and coefficients generated over many decades of RG evolution.

V. OUTLOOK

A salient feature of flavored accelerator observables is the coexistence of strong and weak contributions to the flavor-changing amplitude, with the weak piece parametrically suppressed. This implies two conceptually distinct

constraints on the PQ scale, the strongest constraint arising in the generic regime of strong-amplitude dominance. Conversely, weak dominance implies a weaker limit but requires highly tuned UV coupling configurations. Precisely the difficulty of realizing such configurations renders the corresponding bound valuable as a conservative general lower limit on f_a . Moreover, these constraints are obtained in a fully controlled laboratory environment and thus complementary to astrophysical constraints.

The analysis described above can be extended to other kaon decays, notably modes with a pion pair plus missing energy, which probe the complementary $(k_A)_{\text{sd}}$ coupling. In the $K^+ \rightarrow \pi^+ \pi^0 a$ channel, NA62 is expected to reach $\mathcal{B} \sim 10^{-7}$ [54], 2 orders of magnitude below the current best limit [55]. Relative to the one-pion-plus-invisible case (Fig. 1), the additional π^0 reduces sensitivity through the smaller trigger bandwidth (by $\sim 1/400$), the geometric acceptance of the two photons in the electromagnetic calorimeter ($\sim 1/4$), and the larger background in three- versus two-body kinematics ($\sim 1/10$), while reconstruction efficiencies for π^+ and π^0 remain comparable. Altogether, this implies that $\mathcal{B}(K^+ \rightarrow \pi^+ \pi^0 a) \lesssim 5 \times 10^{-7}$. For the analogous neutral channel, E391a has set $\mathcal{B}(K_L \rightarrow \pi^0 \pi^0 a) < 7 \times 10^{-7}$ at 90% CL [56], and future improvements are expected from KOTO and KOTO-II [57]. A sensitivity at the level of $\mathcal{B}(K \rightarrow \pi \pi a) \lesssim 5 \times 10^{-7}$ therefore provides a realistic benchmark, translating into

$$|(F_A)_{\text{sd}}| > 1.0 \times 10^8 \text{ GeV}, \quad (10)$$

which exceeds the latest bound in Ref. [51] by one order of magnitude. As for the vector case, this constraint is noteworthy both for its strength and for being derived in a controlled experimental setting.

ACKNOWLEDGMENTS

We acknowledge discussions with Maël Cavan-Piton in the initial stages of the project. A special thanks goes to Claudia Cornella for key clarifications and insights on G_8^θ . This work has received funding from the French National Research Agency (ANR), under Contracts No. ANR-19-CE31-0016 (“GammaRare”) and No. ANR-23-CE31-0018 (“InvISYble”), that we gratefully acknowledge.

DATA AVAILABILITY

The data that support the findings of this article are not publicly available. The data are available from the authors upon reasonable request.

APPENDIX A: VALIDATIONS OF OUR REINTERPRETATION PROCEDURE

We tested the procedure described in the main text by applying it to 2017 [40] and 2016–2018 data [36], and verifying that the obtained upper limits on $\mathcal{B}(K^+ \rightarrow \pi^+ a)$ reproduce the published counterparts in Refs. [36,40].

In Fig. 2, we show our results obtained for the accessible values of axion masses compared to the published ones from the collaboration. We find a very good agreement, to within a few percent, if we compare to the χ^2 value with 90% probability and 2 d.o.f., which we understand to be due to neglecting the uncertainty of the knowledge of the SES.




FIG. 2. Upper limit on $\mathcal{B}(K^+ \rightarrow \pi^+ a)$ from the 2017 [40] and 2016–18 [36] datasets. The solid red line is the published result of NA62 collaboration. The gray dashed line shows the corresponding value of $\delta\mathcal{B}_{\nu\bar{\nu}}$ for comparison.

APPENDIX B: ChPT CONVENTIONS AND $K^+ \rightarrow \pi^+ a$ AMPLITUDES

We collect the Lagrangians relevant to our calculations. The “strong” Lagrangian reads

$$\mathcal{L}_s = \frac{F_0^2}{4} (\text{Tr} D_\mu U (D^\mu U)^\dagger + 2B_0 \text{Tr} (\hat{M}_q U + U^\dagger \hat{M}_q^\dagger)). \quad (\text{B1})$$

The pion field $U \equiv \exp(i\frac{\phi^a}{F_0}\lambda^a)$ is normalized such that $(\phi^a \lambda^a)_{11} = \pi^0 + \eta_8/\sqrt{3}$, with $F_0 = 93$ MeV, and transforms as $U \rightarrow RUL^\dagger$; $D_\mu U = \partial_\mu U - ir_\mu U + iUl_\mu$, with r_μ, l_μ external currents, e.g.,

$$\begin{aligned} r_\mu^{\text{e.m.}} &= l_\mu^{\text{e.m.}} = -eA_\mu Q, \\ l_\mu^W &= -\frac{g}{\sqrt{2}} W_\mu^+ \begin{pmatrix} 0 & V_{ud} & V_{us} \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} + \text{H.c.}, \\ r_\mu^W &= 0, \\ r_\mu^Z &= -\frac{g}{\cos \theta_w} Z_\mu^0 (-Q \sin^2 \theta_w), \\ l_\mu^Z &= -\frac{g}{\cos \theta_w} Z_\mu^0 (T_3 - Q \sin^2 \theta_w), \end{aligned} \quad (\text{B2})$$

with $Q = \text{diag}(Q_u, Q_d, Q_s)$ and $T_3 = \text{diag}(1/2, -1/2, -1/2)$.

As elucidated in Ref. [25], the above Lagrangian assumes that the $aG\tilde{G}$ coupling in Eq. (1) has been removed in favor of an axion-dependent phase in the quark-mass matrix $M_q = \text{diag}(m_u, m_d, m_s)$. We define this quark redefinition as $q \rightarrow \exp[i(\delta + \kappa\gamma_5)a/(2f_a)]q$, with $\text{Tr}\kappa = 1$ and the Hermitian matrices δ, κ chosen as diagonal, implying that [25,27]

$$\begin{aligned} \hat{M}_q &= \exp\left(i\frac{a}{f_a}\kappa\right)M_q, \\ r_\mu^a &= \frac{\partial_\mu a}{f_a} \hat{k}_R(a), \\ l_\mu^a &= \frac{\partial_\mu a}{f_a} \hat{k}_L(a), \end{aligned} \quad (\text{B3})$$

with $\hat{k}_{R,L}(a) = e^{-ia/(2f_a)(\delta \pm \kappa)}(k_{R,L} - (\delta \pm \kappa)/2)e^{ia/(2f_a)(\delta \pm \kappa)}$. Note that the κ and δ dependence cancels out in physical amplitudes such as Eqs. (B7) and (B8), as expected. For $G\tilde{G}$, we use the normalization $\tilde{G}_a^{\mu\nu} \equiv \frac{1}{2}\epsilon^{\mu\nu\rho\sigma}G_{\rho\sigma,a}$ with $\epsilon^{0123} = +1$ as in Ref. [59]. Our effective scale f_a definition thus coincides with the one in Refs. [9,25].

The $|\Delta S| = 1$ Lagrangian is given by [43]

$$\mathcal{L}_w = 4G_8(L_\mu L^\mu)^{32} + 2F_0^2 G_8^\theta (D_\mu \theta)(L^\mu)^{32} + \text{H.c.}, \quad (\text{B4})$$

with $G_8^{(\theta)} \simeq -\frac{G_F}{\sqrt{2}} V_{ud}^* V_{us} g_8^{(\theta)}$ and with the L_μ current

$$L_\mu^{ji} = i\frac{F_0^2}{2} \exp -ia/(2f_a) ((\delta_i - \delta_j) - (\kappa_i - \kappa_j)) [(D_\mu U)^\dagger U]^{ji}. \quad (\text{B5})$$

The covariant derivative of the external current θ has been introduced as in Ref. [43] and is defined as

$$D_\mu \theta \equiv \partial_\mu \theta - \langle r_\mu - l_\mu \rangle = -\langle r_\mu - l_\mu \rangle, \quad (\text{B6})$$

where the last equality comes from the fact that the axion-dependent quark redefinition introduced above set $\theta = 0$. As discussed in the main text, we find that the couplings in the two operators in Eq. (B4) are not independent, as they fulfill the relation $G_8^\theta = -2G_8$, or equivalently $g_8^\theta = -2g_8$.

Throughout, we mostly follow the conventions in Refs. [27,58], with the exception of the convention on the chiral transformation properties of the meson field, the pion decay constant, and the g_8 definition, for which we comply with classical ChPT papers such as Ref. [60]. Numerically, we take $g_8 = 3.61$ [53].

For the $K^+ \rightarrow \pi^+ a$ amplitudes, we find that

$$i\mathcal{M}_s = \frac{m_K^2(k_V)_{sd}}{2f_a}(1 - x_{\pi K}), \quad (\text{B7})$$

$$\begin{aligned} i\mathcal{M}_w = & -\frac{N_8^*}{4f_a} \left(\frac{m_K^4 C_K + m_K^2 m_\pi^2 C_{K\pi} + m_\pi^4 C_\pi + m_a^2 m_K^2 C_{aK} + m_a^2 m_\pi^2 C_{a\pi} + m_a^4 C_a}{4m_K^2 - m_\pi^2 - 3m_a^2} \right) \\ & - \frac{N_8^{\theta*}}{2f_a} (m_K^2 - m_\pi^2) (-1 + (k_A)_{uu} + (k_A)_{dd} + (k_A)_{ss}) \\ = & -\frac{N_8^* m_K^2}{4f_a} \left(-2 + 2(k_A)_{uu} + (k_A)_{dd} + (k_A)_{ss} + (k_V)_{dd} - (k_V)_{ss} + 2 \frac{N_8^{\theta*}}{N_8^*} (-1 + (k_A)_{uu} + (k_A)_{dd} + (k_A)_{ss}) \right. \\ & \left. + O(x_{aK}, x_{\pi K}) \right), \quad (\text{B8}) \end{aligned}$$

with $x_{ij} \equiv (m_i/m_j)^2$ and $N_8^{(\theta)} \equiv 2F_0^2 G_8^{(\theta)}$. In the first line for $i\mathcal{M}_w$, contributions are ordered by coupling (N_8 versus N_8^θ) and by expected numerical significance in the QCD-axion case ($x_{a\pi} \ll 1$), and bearing in mind that $x_{\pi K} \ll 1$. The last line is an explicit expansion in these small ratios. The C symbols encode the dependence on the fundamental axion-quark couplings as

$$\begin{aligned} C_K &= 4(-2 + 2(k_A)_{uu} + (k_A)_{dd} + (k_A)_{ss} + (k_V)_{dd} - (k_V)_{ss}), \\ C_{K\pi} &= 8 - 10(k_A)_{uu} - 5(k_A)_{dd} - 5(k_A)_{ss} + 3(k_V)_{dd} - 3(k_V)_{ss}, \\ C_\pi &= 2(k_A)_{uu} + (k_A)_{dd} + (k_A)_{ss} - (k_V)_{dd} + (k_V)_{ss}, \\ C_{aK} &= 8 - 8(k_A)_{uu} - (k_A)_{dd} - 3(k_A)_{ss} - 7(k_V)_{dd} + 7(k_V)_{ss}, \\ C_{a\pi} &= 2(-4 + 4(k_A)_{uu} + 2(k_A)_{dd} - (k_V)_{dd} + (k_V)_{ss}), \\ C_a &= -3((k_A)_{dd} - (k_A)_{ss} - (k_V)_{dd} + (k_V)_{ss}). \quad (\text{B9}) \end{aligned}$$

In the above equations, the k couplings are understood to be at a renormalization scale μ close to the scale of the process, i.e., at $\mu_K \approx m_K$.

APPENDIX C: MASS-COUPLING RELATION

In QCD-axion models, the mass can be related to the effective flavor-violating scale bounded in this work through

$$m_a = \frac{\sqrt{2}F_\pi m_\pi}{|(F_V)_{sd}| |(k_V)_{sd}|} \frac{\sqrt{m_u m_d}}{m_u + m_d}, \quad (\text{C1})$$

with $F_\pi = 93\sqrt{2}$ MeV. This relation allows to display the impact of the $(F_V)_{sd}$ bound [Eq. (6)] in the mass-versus-coupling plane. This is shown in Fig. 3, where we restrict to $|(k_V)_{sd}| \lesssim 10$ to remain within the perturbative regime.

APPENDIX D: f_a LIMIT AS A FUNCTION OF THE UV COUPLINGS

1. Linear form and hierarchy of coefficients

The bound on the axion decay constant derived in the main text can be written in the compact form

$$f_a \geq \tilde{f}_a |z|, \quad z \equiv 1 + \alpha \cdot \mathbf{c}, \quad (\text{D1})$$

where $\mathbf{c} \equiv (c_1, \dots, c_{35})$ denotes the set of 35 dimensionless axion couplings defined at the PQ scale μ_{PQ} , and the numerical coefficients α_i subsume both strong and weak contributions to the $K \rightarrow \pi a$ amplitude, and include renormalization-group evolution to the scale μ_K of the decay.

The full list of coefficients α_i , evaluated at the reference scale $\mu_{\text{PQ}} = 10^{12}$ GeV, is reported in Table I. A salient feature is the very wide hierarchy in their magnitudes: $|\alpha_i| \in [O(10^{-2}), O(10^6)]$. This hierarchy directly reflects the co-existence of two qualitatively different contributions to the decay amplitude: a strong contribution, proportional to the flavor-violating axion coupling $(k_V)_{sd}$, and a weak contribution, which comes with a relative parametric suppression of order $G_F F_0^2 \sim 10^{-7}$. The largest values of $|\alpha_i|$ are associated with couplings that feed the strong amplitude through renormalization-group mixing, while the smallest coefficients correspond to purely weak or electroweak-induced effects.

Equation (D1) makes explicit that the bound on f_a depends linearly on the UV couplings through the complex quantity z . In the absence of special correlations among the components of $\mathbf{c}$, the magnitude of z is therefore expected to be set by the largest coefficients in Table I. The implications of this structure for the typical size of $|z|$, and for the plausibility of accidental cancellations leading to $|z| \sim O(1)$, are analyzed in detail in the following subsections.




FIG. 3. QCD axion mass m_a as a function of the effective flavor-violating coupling $(k_V)_{sd}$.




FIG. 4. Distribution of f_a bounds as defined in Eq. (8), assuming that $c_i \sim \text{Unif}[-5, 5]$: (left) scan of all UV couplings; (right) scan with UV flavor-violating couplings set to zero.

2. Generic behavior and geometric interpretation

We analyze here the generic behavior of z , as defined in Eq. (D1), entering the bound [Eq. (8)], under minimal assumptions on the UV couplings $\mathbf{c}$.

Typical scale of the amplitude. As shown in Table I, the coefficients α_i span several orders of magnitude, up to $|\alpha_i| \sim O(10^6)$. For bounded UV couplings $c_i \in [-5, 5]$ with no special correlations, the linear combination $\alpha \cdot \mathbf{c}$ therefore has a typical magnitude $|1 + \alpha \cdot \mathbf{c}| \sim O(10^7)$, so that $|z|$ is generically dominated by this term rather than by the constant 1. This expectation is confirmed by the numerical distribution shown in Fig. 4 (left).

Induced distribution in the complex plane. Since z is a linear function of the real variables c_i , any choice of a joint distribution for c_i with a smooth density and finite variance induces a smooth probability density $f(z)$ for the complex variable z . No assumption of Gaussianity or independence is required for what follows; only local smoothness of $f(z)$ near the origin will be used.

Small- $|z|$ scaling. The “small-ball” probability for z to lie inside a disk of radius r around the origin is

$$P(|z| \leq r) = \int_{|z| \leq r} f(z) d^2z = \pi r^2 f(0) + o(r^2) \quad (r \rightarrow 0), \quad (\text{D2})$$

where the first equality is by definition and the second follows from a Taylor expansion, assuming only smoothness of $f(z)$ at the origin and otherwise independently of its detailed form. For $r \ll |z|_{\text{peak}}$, the cumulative probability therefore scales as r^2 . This behavior is purely geometric, reflecting the area of a disk in the complex plane and not relying on any further distributional assumption.

Extrapolation from finite scans. The quadratic scaling implies that the probability of reaching the weak-amplitude regime $|z| \sim O(1)$ is suppressed by the square of the ratio between the target radius—unity—and the typical radius— $|z|_{\text{peak}} = O(10^7)$ —of the distribution. This can be estimated empirically by determining $P(|z| \leq r_0)$ for intermediate values $1 \ll r_0 \ll |z|_{\text{peak}}$ and extrapolating according to $P(|z| \leq r) \propto r^2$. Applying this procedure to our numerical scan shows

the onset of a stable plateau for $r_0 \lesssim 10^6$, corresponding to

$$P(|z| \leq 1) \simeq 1 \times 10^{-16}, \quad (\text{D3})$$

as also illustrated quantitatively by Fig. 5 (right) and by the accompanying Table II. Consequently, obtaining one single instance with $|z| \sim O(1)$ would require a scan of order 10^{15} points, far beyond any practical brute-force exploration of the parameter space.

Geometric interpretation. Geometrically, the complex vector $\alpha \cdot \mathbf{c}$ typically populates a region of radius $\sim 10^7$ in the complex plane. Reaching $|z| \sim O(1)$ corresponds to forcing this vector to lie inside a disk whose radius is smaller by 7 orders of magnitude. Such configurations therefore occupy a vanishingly small fraction of the full 35-dimensional parameter space and require highly correlated cancellations among many independent contributions. Two simple examples of such cancellations will be explored in Appendix D3.

These considerations show that the weak-amplitude regime $|z| \sim O(1)$, while not excluded mathematically, is geometrically exceptional. This supports the interpretation of the bound [Eq. (9)] as a conservative absolute limit on f_a :




FIG. 5. Empirical determinations of $P(|z| \leq 1)$ from Monte Carlo evaluations of $P(|z| \leq r_0)$, with r_0 values in the “intermediate” range $1 \ll r_0 \ll |z|_{\text{peak}}$ corresponding to the onset of r^2 scaling. See text and Table II for further details.

TABLE II. Small-ball probabilities from a Monte Carlo scan with $N = 1 \times 10^{10}$ samples and $c_i \sim \text{Unif}[-5.0, 5.0]$. We report the number $N_{|z| \leq r_0}$ of samples satisfying $|z| \leq r_0$, the corresponding empirical probability $P(|z| \leq r_0) = N_{|z| \leq r_0}/N$, and the extrapolated $\hat{P}(|z| \leq 1) = P(|z| \leq r_0)/r_0^2$, with δ denoting the propagated binomial uncertainty.

r_0	$N_{ z \leq r_0}$	$P(z \leq r_0)$	$\hat{P}(z \leq 1) \pm \delta$
1.0×10^4	109	1.090×10^{-8}	$1.09(10) \times 10^{-16}$
2.0×10^4	423	4.230×10^{-8}	$1.058(51) \times 10^{-16}$
5.0×10^4	2708	2.708×10^{-7}	$1.083(21) \times 10^{-16}$
1.0×10^5	10 733	1.073×10^{-6}	$1.073(10) \times 10^{-16}$
5.0×10^5	269 426	2.694×10^{-5}	$1.0777(21) \times 10^{-16}$
1.0×10^6	1 076 190	1.076×10^{-4}	$1.0762(10) \times 10^{-16}$
2.0×10^6	4 286 966	4.287×10^{-4}	$1.07174(52) \times 10^{-16}$
5.0×10^6	26 401 016	2.640×10^{-3}	$1.05604(21) \times 10^{-16}$

In the overwhelming majority of parameter space, the implied constraint is stronger by many orders of magnitude.

3. On specific configurations yielding weak-amplitude dominance

In this Appendix, we examine quantitatively whether the bound [Eq. (9)] can be approached in special corners of parameter space. We analyze two representative situations that can lead to $|1 + \alpha \cdot \mathbf{c}| \sim O(1)$: (1) dominance by a single coupling and (2) engineered cancellations between one coupling and the sum of the rest. We consider them in turn.

a. Case (1): single-coupling dominance

We first consider the situation in which only a single UV coupling c_k is varied, while all other $c_{j \neq k}$ are set to zero. In this case, $|z| = |1 + \alpha_k c_k|$ (no sum over k), and the bound on f_a depends only on $c_k \in [-5, 5]$. Except in a small neighborhood of $c_k \simeq -1/\alpha_k$, one finds that $|1 + \alpha_k c_k|$ is of order unity or larger, so that the resulting bound on f_a is equal to or stronger than $\tilde{f}_a$.

A weaker bound can arise if c_k is tuned to cancel the constant term, namely, $c_k \simeq -1/\alpha_k$, which produces a cusp-shaped minimum in $f_a(c_k)$. This behavior is illustrated in Fig. 6, where the minimum of the bound occurs at $c_k = -1/\alpha_k$ whenever this value lies within $[-5, 5]$. In some cases, the minimum of f_a falls slightly below $\tilde{f}_a$, but only for comparatively large couplings (magnitude ~ 2) and only when the mentioned cancellation condition is enforced.

Such a cancellation requires a specific numerical relation between two quantities that are *a priori* unrelated: the UV coupling c_k , defined at the PQ scale, and the coefficient α_k , which encodes renormalization-group evolution over many decades in energy. As no known symmetry or dynamical argument correlates these quantities, achieving $c_k \simeq -1/\alpha_k$ amounts to fine-tuning. We therefore conclude that single-parameter dominance does not generically evade the bound [Eq. (9)].

b. Case (2): cancellations between one parameter and the sum of the rest

We now consider the possibility that one UV coupling c_i cancels the remaining contributions to $z = 1 + \alpha \cdot \mathbf{c}$, leading

to $|z| \sim O(1)$. This requires

$$c_i \simeq -\frac{1}{\alpha_i} \left(1 + \sum_{j \neq i} \alpha_j c_j \right). \quad (\text{D4})$$

The 35 independent parameters c_i comprise the diagonal gauge and fermion couplings as well as the real and imaginary parts of off-diagonal fermion couplings; in all cases, c_i is real by construction. The right-hand side of Eq. (D4), however, is in general complex. As a result, Eq. (D4) amounts to two independent conditions: One fixes c_i to the real part of the right-hand side, while the other enforces the vanishing of its imaginary part.

To see when the first condition can be satisfied without additional tuning, rewrite the cancellation condition as

$$\alpha_i c_i \simeq -\sum_{j \neq i} \alpha_j c_j.$$

If $|\alpha_i|$ is among the largest coefficients, the ratios $|\alpha_j/\alpha_i|$ are at most $O(1)$, and the equality can be fulfilled with generic values of the coefficients c_j . These parameters represent UV couplings and are therefore expected to remain perturbative, i.e., not parametrically larger than unity. In contrast, if $|\alpha_i|$ is not among the largest, sum necessarily contains terms with $|\alpha_j/\alpha_i| \gg 1$. Under the same perturbativity requirement, the c_j cannot compensate such enhanced contributions individually; satisfying the condition then requires correlated cancellations among several c_j .

Even if the real-part condition is satisfied, the additional requirement that the imaginary part vanish imposes an independent constraint. It demands correlated choices among several *a priori* unrelated complex contributions. These involve unrelated phases and display a wide hierarchy in magnitude, both features originating from the coefficients α_j , which arise from RG evolution over many decades in energy and thereby link IR and UV dynamics. Given the absence of any symmetry relating them, such an imaginary-part cancellation is strongly nongeneric.

These statements are illustrated in Figs. 7 and 8, which display the distributions of the real and imaginary parts, respectively, of the right-hand side of Eq. (D4) for each of the 35 possible choices of c_i . In the plots, the parameters corresponding to the real and imaginary parts of off-diagonal fermion couplings, $\Re(c_{Q,u,d})_{ij}$ and $\Im(c_{Q,u,d})_{ij}$, are abbreviated as $c_{Q,u,d}^{R,I}{}_{ij}$ to avoid clutter. These denote independent real scan parameters; the $\Re$ and $\Im$ operations shown in the figures always refer to the right-hand side of Eq. (D4). The results are obtained from a Monte Carlo scan with 2×10^7 random samples per parameter.

In each panel, the red dashed line indicates the average absolute value of the distribution, which characterizes the typical scale generated by random UV coefficients; its numerical value appears as the first legend entry. The second legend entry reports the smallest absolute value obtained for that parameter as a result of scanning over the remaining 34 couplings. The third legend entry identifies the corresponding α_i coefficient (see Table I), making explicit which parameters are associated with the largest RG-induced weights.




FIG. 6. f_a limit [Eq. (8)] as a function of one single parameter at a time in the range $[-5, 5]$. The inset shows the minimum f_a limit within this range.


 FIG. 7. Distributions of the real part of Eq. (D4), for all possible choices of c_i . See text for details on the legend entries.

FIG. 8. Distributions of the imaginary part of Eq. (D4), for all possible choices of c_i . See text for details on the legend entries.

Only the four parameters corresponding to the largest $|\alpha_i|$ admit support for the real-part condition within the allowed interval $[-5, 5]$. Even in these cases, the imaginary-part distribution remains strongly peaked away from zero. The last legend entry reports the fraction of scan points satisfying both conditions simultaneously within a numerical tolerance $\epsilon = 0.01$. This fraction is found to be zero or statistically negligible in all cases. (The tolerance $\epsilon = 0.01$ is chosen to be

small compared to the typical scale of the distributions; tightening it further only decreases the already negligible fraction of satisfying points within the explored scan size.)

We conclude that multiparameter cancellations capable of yielding $|z| \sim O(1)$ are not excluded mathematically, but are geometrically and statistically implausible. This reinforces the interpretation of the bound [Eq. (9)] as a general conservative lower limit on f_a .

-
- [1] I. G. Irastorza and J. Redondo, New experimental approaches in the search for axion-like particles, *Prog. Part. Nucl. Phys.* **102**, 89 (2018).
 - [2] G. Lanfranchi, M. Pospelov, and P. Schuster, The search for feebly interacting particles, *Annu. Rev. Nucl. Part. Sci.* **71**, 279 (2021).
 - [3] P. Sikivie, Invisible axion search methods, *Rev. Mod. Phys.* **93**, 015004 (2021).
 - [4] G. G. Raffelt, Astrophysical methods to constrain axions and other novel particle phenomena, *Phys. Rep.* **198**, 1 (1990).
 - [5] G. G. Raffelt, Astrophysical axion bounds, *Lect. Notes Phys.* **741**, 51 (2008).
 - [6] A. Caputo and G. Raffelt, Astrophysical axion bounds: The 2024 edition, in 1st training school of the COST action COSMIC WISPerS (CA21106), *Lect. Notes Phys.* **741**, 51 (2008).
 - [7] D. J. E. Marsh, Axion cosmology, *Phys. Rep.* **643**, 1 (2016).
 - [8] J. E. Kim and G. Carosi, Axions and the strong CP problem, *Rev. Mod. Phys.* **82**, 557 (2010); Erratum: **91**, 049902(E) (2019).
 - [9] L. Di Luzio, M. Giannotti, E. Nardi, and L. Visinelli, The landscape of QCD axion models, *Phys. Rep.* **870**, 1 (2020).
 - [10] K. Choi, S. H. Im, and C. Sub Shin, Recent Progress in the Physics of Axions and Axion-Like Particles, *Annu. Rev. Nucl. Part. Sci.* **71**, 225 (2021).
 - [11] J. E. Kim, Weak interaction singlet and strong CP invariance, *Phys. Rev. Lett.* **43**, 103 (1979).
 - [12] M. A. Shifman, A. I. Vainshtein, and V. I. Zakharov, Can confinement ensure natural CP invariance of strong interactions? *Nucl. Phys. B* **166**, 493 (1980).
 - [13] M. Dine, W. Fischler, and M. Srednicki, A simple solution to the strong CP problem with a harmless axion, *Phys. Lett. B* **104**, 199 (1981).
 - [14] A. R. Zhitnitsky, On possible suppression of the axion hadron interactions (in Russian), *Sov. J. Nucl. Phys.* **31**, 260 (1980).
 - [15] R. D. Peccei and H. R. Quinn, CP conservation in the presence of instantons, *Phys. Rev. Lett.* **38**, 1440 (1977).
 - [16] R. D. Peccei and H. R. Quinn, Constraints imposed by CP conservation in the presence of instantons, *Phys. Rev. D* **16**, 1791 (1977).
 - [17] S. Weinberg, A new light boson? *Phys. Rev. Lett.* **40**, 223 (1978).
 - [18] F. Wilczek, Problem of strong P and T invariance in the presence of instantons, *Phys. Rev. Lett.* **40**, 279 (1978).
 - [19] J. Preskill, M. B. Wise, and F. Wilczek, Cosmology of the invisible axion, *Phys. Lett. B* **120**, 127 (1983).
 - [20] L. F. Abbott and P. Sikivie, A cosmological bound on the invisible axion, *Phys. Lett. B* **120**, 133 (1983).
 - [21] M. Dine and W. Fischler, The not so harmless axion, *Phys. Lett. B* **120**, 137 (1983).
 - [22] C. L. Cowan, F. Reines, F. B. Harrison, H. W. Kruse, and A. D. McGuire, Detection of the free neutrino: A confirmation, *Science* **124**, 103 (1956).
 - [23] J. Gasser and H. Leutwyler, Chiral perturbation theory to one loop, *Ann. Phys.* **158**, 142 (1984).
 - [24] S. Weinberg, Phenomenological Lagrangians, *Physica A* **96**, 327 (1979).
 - [25] H. Georgi, D. B. Kaplan, and L. Randall, Manifesting the invisible axion at low-energies, *Phys. Lett. B* **169**, 73 (1986).
 - [26] Our definition of the Peccei-Quinn scale coincides with that in the review [9]. The sign difference in the defining $aG\tilde{G}$ term in Eq. (1) arises from our use of opposite ϵ -tensor conventions relative to those in Ref. [9] (see Appendix B).
 - [27] M. Bauer, M. Neubert, S. Renner, M. Schnubel, and A. Thamm, Consistent treatment of axions in the weak chiral Lagrangian, *Phys. Rev. Lett.* **127**, 081803 (2021).
 - [28] M. Buschmann, C. Dessert, J. W. Foster, A. J. Long, and B. R. Safdi, Upper limit on the QCD axion mass from isolated neutron star cooling, *Phys. Rev. Lett.* **128**, 091102 (2022).
 - [29] M. Cavan-Piton, D. Guadagnoli, M. Oertel, H. Seong, and L. Vittorio, Axion emission from strange matter in core-collapse SNe, *Phys. Rev. Lett.* **133**, 121002 (2024).
 - [30] T. Fischer, J. Martin Camalich, H. Kochankovski, and L. Tolos, Hyperons during proto-neutron star deleptonization and the emission of dark flavored particles, *J. Cosmol. Astropart. Phys.* **01** (2025) 061.
 - [31] A. J. Buras and E. Venturini, The exclusive vision of rare K and B decays and of the quark mixing in the standard model, *Eur. Phys. J. C* **82**, 615 (2022).
 - [32] G. Anzivino *et al.*, Workshop summary: Kaons@CERN 2023, *Eur. Phys. J. C* **84**, 377 (2024).
 - [33] G. D'Ambrosio, A. M. Iyer, F. Mahmoudi, and S. Neshatpour, Anatomy of kaon decays and prospects for lepton flavor universality violation, *J. High Energy Phys.* **09** (2022) 148.
 - [34] E. Cortina Gil *et al.* (NA62 Collaboration), First search for $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ using the decay-in-flight technique, *Phys. Lett. B* **791**, 156 (2019).
 - [35] E. Cortina Gil *et al.* (NA62 Collaboration), An investigation of the very rare $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ decay, *J. High Energy Phys.* **11** (2020) 042.
 - [36] E. Cortina Gil *et al.* (NA62 Collaboration), Measurement of the very rare $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ decay, *J. High Energy Phys.* **06** (2021) 093.
 - [37] E. Cortina Gil *et al.* (NA62 Collaboration), Observation of the $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ decay and measurement of its branching ratio, *J. High Energy Phys.* **02** (2024) 191.

- [38] E. Cortina Gil *et al.* (NA62 Collaboration), Searches for hidden sectors using $K^+ \rightarrow \pi^+ X$ decays, *J. High Energy Phys.* **11** (2025) 143.
- [39] S. Alibocus *et al.* (NA62 Collaboration), Measurement of the branching ratio of the $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ decay, [arXiv:2607.16413](https://arxiv.org/abs/2607.16413).
- [40] E. Cortina Gil *et al.* (NA62 Collaboration), Search for a feebly interacting particle X in the decay $K^+ \rightarrow \pi^+ X$, *J. High Energy Phys.* **03** (2021) 058.
- [41] Alternatively, one may treat $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ events as a separate component n_ν left as a free parameter. This approach will become viable in experimental analyses as the observed n_ν increases. For a fixed number of d.o.f., we find that allowing n_ν to float strengthens the bounds by approximately 5%–10% compared to fixing it to the SM value. However, we adhere to the latter approach to align fully with the procedure outlined in Ref. [40].
- [42] R. D. Cousins, J. T. Linnemann, and J. Tucker, Evaluation of three methods for calculating statistical significance when incorporating a systematic uncertainty into a test of the background-only hypothesis for a poisson process, *Nucl. Instrum. Methods Phys. Res. Sect. A* **595**, 480 (2008).
- [43] C. Cornella, A. M. Galda, M. Neubert, and D. Wyler, $K^\pm \rightarrow \pi^\pm a$ at next-to-leading order in chiral perturbation theory and updated bounds on ALP couplings, *J. High Energy Phys.* **06** (2024) 029.
- [44] G. Isidori, F. Mescia, and C. Smith, Light-quark loops in $K \rightarrow \pi \nu \bar{\nu}$, *Nucl. Phys. B* **718**, 319 (2005).
- [45] S. Di Noi and L. Silvestrini, RGEsolver: A C++ library to perform renormalization group evolution in the standard model effective theory, *Eur. Phys. J. C* **83**, 200 (2023).
- [46] H. Arason, D. J. Castano, B. Keszthelyi, S. Mikaelian, E. J. Piard, P. Ramond, and B. D. Wright, Renormalization group study of the standard model and its extensions. I. The Standard model, *Phys. Rev. D* **46**, 3945 (1992).
- [47] M. Bauer, M. Neubert, S. Renner, M. Schnubel, and A. Thamm, The low-energy effective theory of axions and ALPs, *J. High Energy Phys.* **04** (2021) 063.
- [48] For the QCD axion, the coupling to gluons defines the PQ scale [9] [see Eq. (1)]. For the normalization of the $G\tilde{G}$ operator, see Appendix B.
- [49] In a different basis, off-diagonal entries of the axion-down-type quark coupling matrix are also induced by the misalignment between the gauge and the mass eigenstates; see e.g., the first term on the right side of Eq. (5.17) of Ref. [47] for the case of choosing the up-quark basis. More generally, the k couplings are defined up to the biunitary transformations of the up- and down-type quark fields that diagonalize the Yukawa matrices. Each of these transformations can induce off-diagonalities even starting from diagonal k matrices.
- [50] Specifically, RG effects in off-diagonal k couplings are suppressed below μ_{EW} by the Fermi constant and the tiny Yukawa couplings of light quarks [47,58]; the same effects in diagonal k couplings enter $\mathcal{M}_w$ only, and thus amount to contributions to the amplitude of $O(G_F^2)$.
- [51] Jorge Martin Camalich, M. Pospelov, P. N. H. Vuong, R. Ziegler, and J. Zupan, Quark flavor phenomenology of the QCD axion, *Phys. Rev. D* **102**, 015023 (2020).
- [52] C. O'Hare, cajohare/axionlimits: Axionlimits, 2020, <https://cajohare.github.io/AxionLimits/>.
- [53] V. Cirigliano, G. Ecker, H. Neufeld, A. Pich, and J. Portoles, Kaon decays in the standard model, *Rev. Mod. Phys.* **84**, 399 (2012).
- [54] E. Goudzovski *et al.*, New physics searches at kaon and hyperon factories, *Rep. Prog. Phys.* **86**, 016201 (2023).
- [55] O. G. Tchikilev *et al.*, Search for light pseudoscalar sgoldstino in K^- decays, *Phys. Lett. B* **602**, 149 (2004).
- [56] R. Ogata *et al.* (E391a Collaboration), Study of the $K_L^0 \rightarrow \pi^0 \pi^0 \nu \bar{\nu}$ decay, *Phys. Rev. D* **84**, 052009 (2011).
- [57] J. Fry *et al.* (KOTO Collaboration), Proposal of the KOTO II experiment, [arXiv:2501.14827](https://arxiv.org/abs/2501.14827).
- [58] M. Bauer, M. Neubert, S. Renner, M. Schnubel, and A. Thamm, Flavor probes of axion-like particles, *J. High Energy Phys.* **09** (2022) 056.
- [59] M. E. Peskin and D. V. Schroeder, *An Introduction to Quantum Field Theory* (Addison-Wesley, Reading, 1995).
- [60] G. Ecker, A. Pich, and E. de Rafael, $K^+ \rightarrow \pi^0 \ell^+ \nu_\ell$ decays in the effective chiral Lagrangian of the standard model, *Nucl. Phys. B* **291**, 692 (1987).