Dissertation

submitted to the
Combined Faculty of Natural Sciences and Mathematics
of the
Ruprecht-Karls-University Heidelberg
for the degree of
Doctor of Natural Sciences

Put forward by

Lennert Jarl Thormählen

born in Hamburg

Disputation on 27 October 2022

Linking QCD axion models to their low-energy phenomenology

INSTITUTE FOR THEORETICAL PHYSICS
DEPARTMENT OF PHYSICS AND ASTRONOMY
RUPRECHT-KARLS-UNIVERSITY HEIDELBERG

Referees: **Prof. Dr. Jörg Jäckel**
Prof. Dr. Jan M. Pawlowski

Linking QCD axion models to their low-energy phenomenology

The QCD axion is a well-motivated hypothetical particle beyond the Standard Model of particle physics. It can solve the strong CP problem and account for dark matter in the universe. An axion is a pseudo-Goldstone boson of a spontaneously broken global U(1) symmetry with a chiral anomaly in the strong sector. Such a minimal construction allows for a large number of alternative realisations with similar phenomenology. In this thesis, links between the underlying high-energy physics of axion models and the resulting observable phenomena at low energy are investigated. First, the property of periodicity along with its realisation and consequences at lower energies are studied. Subsequently, the usual leading logarithm approximation for the calculation of loop-induced flavour-changing effects is critically examined. A new model is developed that leads to large logarithmic enhancements and is valid up to the scale of spontaneous symmetry breaking. The phenomenology of this model in future searches for solar axions and at colliders is discussed. Finally, the spectrum of solar axions is evaluated, improving on the previous calculation and systematically quantifying the associated uncertainties for the first time. This could enable the distinction of axion models as well as the measurement of unknown solar quantities with the upcoming helioscope IAXO. It is also shown that this experiment is sensitive to a singular line in the spectrum originating from axion-nucleon interactions.

Verbindungen zwischen QCD Axion Modellen und ihrer Phänomenologie bei niedriger Energie

Das QCD Axion ist ein gut motiviertes hypothetisches Teilchen jenseits des Standardmodells der Teilchenphysik. Es kann das starke CP-Problem lösen und die dunkle Materie im Universum ausmachen. Ein Axion ist ein Pseudo-Goldstone-Boson einer spontan gebrochenen globalen U(1)-Symmetrie mit einer chiralen Anomalie im starken Sektor. Eine solche minimale Konstruktion ermöglicht eine große Anzahl von alternativen Realisierungen mit ähnlicher Phänomenologie. In dieser Arbeit werden Zusammenhänge zwischen der zugrundeliegenden Hochenergiephysik von Axionmodellen und den daraus resultierenden beobachtbaren Phänomenen bei niedriger Energie untersucht. Dies beginnt bei der Eigenschaft der Periodizität sowie deren Realisierung und Konsequenzen bei kleineren Energien. Anschließend wird die übliche Näherung des führenden Logarithmus für die Berechnung von schleifeninduzierten und flavour-verändernden Effekten kritisch untersucht. Es wird ein neues Modell entwickelt, das zu großen logarithmischen Verstärkungen führt und bis zur Skala der spontanen Symmetriebrechung gültig ist. Die Phänomenologie dieses Modells an Teilchenbeschleunigern und bei zukünftigen Suchen nach solaren Axionen wird diskutiert. Schließlich wird das Spektrum der solaren Axionen bestimmt, wobei die bisherige Berechnung verbessert und die damit verbundenen Unsicherheiten erstmals systematisch quantifiziert werden. Dies könnte die Unterscheidung von Axionmodellen sowie die Messung unbekannter solarer Größen mit dem kommenden Helioskop IAXO ermöglichen. Es wird auch gezeigt, dass dieses Experiment für eine singuläre Linie im Spektrum empfindlich ist, die von Axion-Nukleon-Wechselwirkungen herrührt.

Publications

This thesis contains results from collaborative research, which has already been published at the time of completion of this thesis. The jointly authored publications and my respective contributions are:

[1] G. Alonso-Álvarez, F. Ertas, J. Jaeckel, F. Kahlhoefer & L.J. Thormaehlen, *Hidden photon dark matter in the light of XENON1T and stellar cooling*, *Journal of Cosmology and Astroparticle Physics* **11** (2020) 029, [2006.11243].

I contributed to the initial idea of the paper at the time of the publication of the XENON1T anomaly and wrote the summary of astrophysical bounds and hints, which is used in section 4.4 of this thesis.

[2] G. Alonso-Álvarez, F. Ertas, J. Jaeckel, F. Kahlhoefer & L.J. Thormaehlen, *Leading logs in QCD axion effective field theory*, *Journal of High Energy Physics* **07** (2021) 059, [2101.03173].

I provided several ideas and contributions over the course of the project, performed the mass diagonalization and the loop calculation in the new model. I also contributed to the renormalization discussion and the phenomenological study, in particular regarding IAXO, and was strongly involved in writing the publication. Results of this work are included in chapter 6 and appendices A to D.

[3] S. Hoof, J. Jaeckel & L.J. Thormaehlen, *Quantifying uncertainties in the solar axion flux and their impact on determining axion model parameters*, *Journal of Cosmology and Astroparticle Physics* **09** (2021) 006, [2101.08789].

I contributed major parts to the development of the C++ library, performed the majority of uncertainty calculations or estimates, came up with the calculation beyond the static limit and was strongly involved in writing the publication. Results of this work are included in chapter 7 and appendix E.

[4] L. Di Luzio, J. Galan, M. Giannotti, I.G. Irastorza, J. Jaeckel, A. Lindner, J. Ruz, U. Schneekloth, L. Sohl, L.J. Thormaehlen & J.K. Vogel *Probing the axion-nucleon coupling with the next generation of axion helioscopes*, *European Physical Journal C* **82** (2022) 2, 120, [2111.06407].

I performed all of the sensitivity calculations, searched for possible alternatives to ^{57}Fe and provided the full initial draft of this work, which makes up the majority of the final publication. Results of this work are included in section 7.5 and appendix F.

Contents

Appendix

1 Introduction

One of the great success stories of modern physics is the development of quantum field theories (QFTs) as a generic language to describe models in particle physics [5]. Formal developments of these theories were simultaneously accompanied by experimental progress and discoveries mostly using particle accelerators. In combination, these efforts lead to the development of the Standard Model (SM) of particle physics [6–9], which is the most successful QFT today in describing the fundamental interactions of nature. A large number of laboratory experiments as well as astrophysical and cosmological observations are in agreement with the SM at an impressive level of accuracy. The pinnacle of this success was the discovery of the Higgs boson at the Large Hadron Collider (LHC) [10, 11] in 2012, which was the final particle required for describing the SM as a consistent gauge theory.

Despite these huge leaps in understanding the fundamental building blocks of nature, the end of physics is very far from being reached. A large number of questions remains unanswered. It is unknown what dark matter, which makes up the majority of matter in the universe, is made of. The mechanism of baryogenesis, the overproduction of matter in comparison to antimatter, is not understood and it is unclear how the non-vanishing neutrino masses are generated. In addition, there are fine-tuning problems, which are at least aesthetic shortcomings but most likely hint towards a lack in understanding the origin of the SM. These are for instance the hierarchy, strong CP and cosmological constant problems. This incomplete list of some of the most pressing issues demonstrates, on the one hand, that physics beyond the SM (BSM) must exist and, on the other hand, that today the most significant discrepancies between observation and theory stem from cosmological observations or mathematical principles rather than laboratory experiments. Furthermore, the large number of issues, which need to be addressed simultaneously, motivates to search for connections among them and a BSM theory successfully solving more than one at a time is in principle favoured over a solution which is fine-tuned to a specific problem.

An example of such a proposed BSM particle, which stands at the centre of this thesis, is the famous QCD axion. It was originally proposed independently by Wilczek [12] and Weinberg [13] as a necessary consequence of the Peccei–Quinn mechanism [14], which solves the strong CP problem. This is the question why the discrete CP symmetry is realised to such a high precision in quantum chromodynamics (QCD). In addition, it was soon realised that the axion has all the properties of a dark matter particle and that it could constitute the dominant type of matter in the universe [15, 16]. In this way the axion addresses two fundamental problems of the SM, making it a particularly attractive BSM candidate. The axion can possibly also be linked to various other problems albeit this may require additional model building. It has for instance been suggested that it could be connected to baryogenesis [17], cosmic inflation [18–23], neutrino masses [24–26] or dark energy [27].

The search strategies required for detecting an axion are also particularly interesting because it lies at a different frontier than the one being explored by large-energy collider experiments. In general, any BSM particle must have properties that have enabled it to escape detection thus far. This can either be because it is so heavy that it cannot be produced with the energy available in modern day accelerators – the high-energy frontier – or because it is so weakly interacting that its influences on experimental setups are too tiny to be detected – the low-energy or precision frontier [28]. These two directions at which new physical phenomena could be hidden are complementary and require vastly different experimental techniques. Because the axion is by construction light and weakly interacting, it lies beyond the current precision frontier, where measurements with an unprecedented accuracy would be required for detection and characterisation of new particles. In recent years and with the LHC getting closer to its maximal potential, the particle physics community has already started to engage more in precision experiments. Specifically for axions, a plethora of new designs have been suggested or are already under construction and the decades to come may indeed bring an axion rush in particle physics.

Besides detecting the signal of a new particle, it is always an important goal of theoretical physics to embed the newly observed phenomena into a consistent theory at all energy scales, i.e. an ultraviolet (UV)-complete model. However, the QCD axion or other similar light pseudoscalars called axion-like particles (ALPs) are a rather generic implication of a large variety of BSM models, particularly in those which feature spontaneous breaking of global symmetries. Distinguishing these models by their low-energy phenomenology is difficult and in some cases it may be impossible. This can be seen as either a blessing or a curse. It means that axion experiments are sensitive to a number of different models, but at the same time, it is hard to pin down or merely retrieve some information about the underlying theory at high energies. For this reason, the thesis at hand works towards a better understanding of links between UV-complete QCD axion models and their low-energy phenomenology. Establishing such links is far from trivial. The reason is that it involves physics at various energy scales as well as methods ranging from theoretical particles physics over solar astrophysics to novel x-ray detection techniques. All of these research fields are crucial for at least some of the results presented in this thesis.

The power of this work becomes clear, when defining what is meant by low-energy phenomenology and comparing this to the scales involved in QCD axion models. In this thesis, observable phenomena are referred to as low-energy when they only contain the axion as an additional degree of freedom to the Standard Model. All of the UV structure of the underlying model, like for instance additional complex scalars, heavy fermions or multiple Higgs fields, are integrated out at this level and one is left with a light pseudoscalar and a set of effective coupling parameters. The energy where this approach is applicable ranges from essentially zero, as in dark matter experiments or light-shining-through wall setups, over keV temperatures inside of stellar objects, all the way to MeV scales in fixed target colliders. Linking such low-energy observables to QCD axion models can eventually allow to make predictions about the physics at the electroweak scale or even significantly higher scales, which are far from reach of present-day or near future collider experiments.

One may wonder whether such an in-depth investigation of what can be learned from an axion signal is called for, given that no axion has been observed so far. The most immediate motivation comes from the large number of aforementioned axion experiments which have been proposed or are already under construction. Of course, it is of immense interest to explore the potential of such searches in order to sensibly set priorities. Also on a more practical level, studies of beyond detection measurements like the ones in this thesis can directly influence the preferred experimental configuration. Furthermore, links between low-energy phenomenology and underlying UV-complete models are of general theoretical interest and may be useful in other contexts than the axion. The lessons learned from this research are therefore applicable regardless of the existence of a QCD axion or the nature of dark matter in our universe.

This thesis is structured as follows. After introducing the axion as a solution to the strong CP problem and a potential DM candidate in chapter 2, the most common axion models are introduced in chapter 3. A brief overview of the most prominent phenomenological implications of axions is provided in chapter 4. This sets the stage for discussing several links between UV models and low-energy phenomena. Chapter 5 is dedicated to an axion property originating from the underlying microscopic theory, namely its periodicity. Even though it is often neglected, it can have significant consequences for setting up effective field theories and can help to review the validity of proposed axion constructions. It is shown how the periodicity is preserved also by loop corrections, which is then used to rule out a proposition connecting the QCD axion to the majoron from neutrino mass generation. Another context in which the intricate connections between low-energy observations and axion models can be studied is provided by UV-divergent observables in axion effective field theories. An example are loop-induced rare meson decays, which are the topic of chapter 6. This eventually leads to the construction of a new QCD axion model for flavour-violating interactions. Finally, the solar axion flux and its uncertainties are thoroughly investigated in chapter 7. This is motivated by upcoming solar axion searches as well as the abundant information content of the solar axion spectrum. Interesting new links are established, like the connection between quantised axion couplings at low energies and solar properties. This allows to disentangle the influences of axion and solar models in section 7.4 and it is an illustrative example of how UV properties like the axion's periodicity can be useful for the interpretation of axion observations. General conclusions, a brief summary of the main results and an outlook on the future of axion physics are given in chapter 8.

2 The QCD axion

The QCD axion is a hypothetical particle which could solve two fundamental problems of the Standard Model of particle physics – the strong CP problem and the dark matter puzzle. This chapter will briefly introduce both problems and explain how they could be addressed by the existence of an axion.

2.1 Theta-term of QCD

Quantum chromodynamics (QCD) is the sector of the SM which describes the strong nuclear interaction. It is defined as an SU(3) gauge theory with a set of six fermions in the fundamental representation called quarks. Its Lagrangian can therefore be written as

$$\mathcal{L}_{\mathrm{QCD}} = -\frac{1}{2}\mathrm{Tr}\, G_{\mu\nu}G^{\mu\nu} + \bar{q}(i\not{D} - M_q)q \ , \tag{2.1}$$

where $G_{\mu\nu}$ is the field strength tensor[1], q are the quark fields, D denotes the gauge covariant derivative and M_q the quark mass matrix. While this description of QCD is sufficient to explain all confirmed interactions between gluons and quarks, it is not complete. Instead, the theta-term, named after its dimensionless prefactor, must be included,

$$\mathcal{L}_\theta = \theta\frac{g_s^2}{16\pi^2}\mathrm{Tr}\, G_{\mu\nu}\widetilde{G}^{\mu\nu} = \theta\frac{\alpha_s}{4\pi}\mathrm{Tr}\, G_{\mu\nu}\widetilde{G}^{\mu\nu} \ . \tag{2.2}$$

g_s or $\alpha_s = g_s^2/(4\pi)$ quantify the strong coupling constant and $\widetilde{G}_{\mu\nu} = \frac{1}{2}\epsilon_{\mu\nu\rho\sigma}G^{\rho\sigma}$ is the dual field strength tensor. There are two ways to argue why this term is a necessary part of QCD. The first is a field-theoretic argument. Quantum field theories are generally defined by imposing local gauge symmetries, global symmetries and by fixing the particle content. Having done so, the defining Lagrangian is found by adding up all renormalizable hermitian operators which can be constructed from the particle content and are admissible by all symmetries. This is because any such operator cannot be excluded on fundamental grounds and is anyway expected to be generated radiatively at the least. Since both G and $\widetilde{G}$ transform in the adjoint representation of SU(3), it is apparent that the theta-term is invariant under all gauge symmetries of the Standard Model and that it should therefore be included.

The second argument relates to the vacuum structure of QCD. It is presented in several reviews, e.g. refs. [29–31]. The starting point are the solutions to the classical field equations in Euclidean space, which have to solve

$$D_\mu G_{\mu\nu} = 0 \ . \tag{2.3}$$

---

[1] In the following, G is regularly expressed in a generator basis T^a such that $G_{\mu\nu} = G^a_{\mu\nu}T^a$ and the usual normalization, Tr $T^aT^b = \frac{1}{2}\delta^{ab}$, is used.

Gauge field configurations with finite Euclidean action solving these equations are called instantons. They are a general phenomenon of non-abelian gauge theories. The arguably simplest example of these are the BPST instantons [32], which are (anti-)self-dual, spherically symmetric and minimise the action. They exist in SU(2) gauge theory and can therefore be embedded in all other SU(N) theories because these contain SU(2) as a subgroup. The self-dual BPST instanton can be expressed in regular gauge as [31]

$$G^a_{\mu\nu}(x; x_0, \rho) = -\frac{4}{g_s}\eta^a_{\mu\nu}\frac{\rho^2}{\left((x-x_0)^2+\rho^2\right)^2} \,. \tag{2.4}$$

The free parameters x_0 and ρ define the position and size of the instanton, respectively. Furthermore, the 't Hooft-symbol $\eta^a_{\mu\nu}$ was introduced, which mixes SU(2) gauge indices ($a = 1, 2, 3$) with Euclidean Lorentz indices ($\mu, \nu = 1, 2, 3, 4$).

$$\eta^a_{\mu\nu} = \begin{cases} \epsilon^{a\mu\nu} & \mu,\nu = 1,2,3 \\ -\delta^{a\nu} & \mu = 4 \\ \delta^{a\mu} & \nu = 4 \\ 0 & \mu = \nu = 4 \end{cases} \tag{2.5}$$

The Euclidean action of this solution is given by $S_\mathrm{E} = 8\pi^2/g_s^2$. This solution demonstrates some of the interesting properties of instantons. At large distances, the field strength decays as $1/(x-x_0)^4$, which means that it will give a finite contribution to the action and that instantons are localised both in Euclidean time and space. Even though the instanton approaches a pure gauge configuration at large distances, the study of homotopy groups shows that the full instanton configuration cannot be continuously deformed into the vacuum solution $G^a_{\mu\nu} = 0$ because the instanton belongs to a different homotopy class. In fact, all field strength configurations which approach pure gauge at infinity can be classified by their homotopy class. These are labelled by the integer Pontryagin index q also called the winding number. It is defined as

$$q = \int \mathrm{d}^4x_\mathrm{E}\frac{\alpha_s}{4\pi}\mathrm{Tr}\, G_{\mu\nu}\widetilde{G}_{\mu\nu} \in \mathbb{Z} \,. \tag{2.6}$$

The BPST instanton defined above gives $q = 1$, while the anti-instanton has $q = -1$. Importantly, q is a topological and gauge invariant and hence a gauge transformation can never map a field configuration from one homotopy class to another.

One may wonder what these classical solutions of the Euclidean field theory have to do with the physical world in Minkowski space. To answer this, first note that one can also define a winding number Q in Minkowski space [31, 33],

$$Q = \int_{\sigma_0}^{\sigma_1} \mathrm{d}^4x\frac{\alpha_s}{4\pi}\mathrm{Tr}\, G_{\mu\nu}\widetilde{G}^{\mu\nu} \,, \tag{2.7}$$

where the integral runs over the volume bounded by two three-dimensional spatial hypersurfaces σ_0 and σ_1 at times t_0 and t_1. Since the integrand is a total derivative (see section 2.2), it can be reduced to two surface integrals, which are integer valued if one assumes a vacuum at the boundaries. While each of these surface integrals is not gauge independent individually, the difference Q is. At a fixed time t, one can therefore not distinguish between the different vacua. However, it does make sense to speak of a *change in the winding number* between times t_0 and t_1 if $Q \neq 0$. One can show that there are no classical solutions to the equations of motion in Minkowski space with $Q \neq 0$ [33, 34]. This is interpreted as the vacuum not being able to change from one topological sector to another in the classical field theory (i.e. the two boundary terms in equation (2.7) are always the same). In a quantum field theory, however, instantons play a crucial role. This is because the Wick rotated instantons, while not being solutions of the equations of motion in Minkowski space, correspond to complex critical points in the path integral and as such are the most probable tunnelling paths from one topological sector to another [33, 34]. Instantons in Minkowski space should therefore be viewed as quantum *processes* (not particles), which can mediate between physically equivalent vacua.[2]

Several lessons can be taken from this discussion of instantons and fixed time winding numbers. First, there is an infinite set of equivalent vacua, denoted $|n\rangle$. Because tunnelling processes exist in the form of instantons, the physical vacuum $|\text{vac}\rangle$ must be a linear superposition of all $|n\rangle$,

$$|\text{vac}\rangle = \sum_{n\in\mathbb{Z}} a_n \, |n\rangle \ . \tag{2.8}$$

Second, because the vacua are physically equivalent, one expects that the operator T which is defined by $T\,|n\rangle = |n+1\rangle$ commutes with the Hamiltonian and that the physical vacuum should only change by a global phase under T [31, 34],

$$\begin{aligned} & T\,|\text{vac}\rangle = e^{-i\theta}\,|\text{vac}\rangle && (2.9)\\ \Rightarrow \quad & |\text{vac}\rangle = \sum_{n\in\mathbb{Z}} e^{in\theta}\,|n\rangle \equiv |\theta\rangle \ . && (2.10)\end{aligned}$$

These are the famous θ-vacua of QCD. They are physically distinguishable states parametrised by an angle θ, which is a free fundamental parameter of the theory. For practical purposes, it is convenient to shift θ from the definition of the physical vacuum to the Lagrangian. Instead of summing over all $|n\rangle$, including the correct phase and restricting the path integral over gauge fields to a fixed topological sector at a time, one can work with the convenient vacuum $|0\rangle$ and perform the full path integral over gauge fields if at the same time one includes the theta-term in the Lagrangian. This way the path integral receives a phase $e^{in\theta}$ for every field configuration with winding number n. In summary, QCD has another free angular parameter θ, which does not appear explicitly in the Lagrangian (2.1). It can either be included as a parameter of the physical vacuum or as the prefactor of the theta-term as in equation (2.2). In this thesis, the latter approach is followed.

---

[2] Instantons also have other noteworthy effects. For instance they induce the determinental 't Hooft interaction [35], which is baryon number violating in the SM electroweak theory.

2.2 The strong CP problem

After recognising the fact that the theta-term is a necessary part of QCD, one needs to investigate its properties and physical consequences. It was already mentioned and used above that it is a total derivative. This can be made explicit by writing $G_{\mu\nu}\widetilde{G}^{\mu\nu}$ as the total derivative of the Chern–Simons current,

$$\mathrm{Tr}\, G_{\mu\nu}\widetilde{G}^{\mu\nu} = \partial^\mu \left(2\epsilon_{\mu\nu\rho\sigma} \mathrm{Tr} \left[A^\nu \partial^\rho A^\sigma - \frac{2}{3} g_s A^\nu A^\rho A^\sigma \right] \right) \ . \tag{2.11}$$

Such a term can only add boundary contributions to the action. This means it will not appear in perturbation theory or Feynman diagrams because all functional derivatives of the action are unaffected by the addition of surface terms. Any physical consequences of the theta-term are therefore caused by non-perturbative effects like instantons.

Notably, the theta-term also violates the discrete CP symmetry, i.e. the combination of charge conjugation and parity inversion.[3] Its transformation properties under C and P can be most intuitively understood by analogy to the electromagnetic theta-term $F_{\mu\nu}\widetilde{F}^{\mu\nu}$, which can be expressed in terms of electric and magnetic fields $\vec{E}$ and $\vec{B}$,

$$F_{\mu\nu}\widetilde{F}^{\mu\nu} \propto \vec{E} \cdot \vec{B} \quad \xrightarrow{\mathrm{CP}} \quad \vec{E} \cdot (-\vec{B}) \propto -F_{\mu\nu}\widetilde{F}^{\mu\nu} \ . \tag{2.12}$$

The gluon field strength tensor transforms precisely in the same way. Therefore, QCD violates CP unless $\theta = 0$ or $\theta = \pi$.[4] In order to measure the value of θ in nature, one needs to probe CP violation related to the strong interaction. The most sensitive observable of this kind is the electric dipole moment of the neutron (nEDM). Because of the absence of a monopole, even a tiny dipole moment could in principle be observed, but so far all measurements are in agreement with no nEDM at all. The strongest upper bound is given by [37, 38]

$$d_n < 1.8 \times 10^{-13}\, e \text{ fm} \ , \tag{2.13}$$

where e denotes the elementary charge. On the theory side, the expected nEDM can be calculated using various techniques like QCD sum rules [39], empirical neutron models [40], chiral perturbation theory [41] or lattice QCD [42–45]. None of these can provide a precise result, but a very conservative estimate of the expected nEDM would be [30]

$$1.2 \times 10^{-3}\, \theta\ e \text{ fm} < d_n < 11 \times 10^{-3}\, \theta\ e \text{ fm} \ . \tag{2.14}$$

This means that an upper bound on the theta-term in nature is given by[5]

$$\theta < 1.5 \times 10^{-10}. \tag{2.15}$$

---

[3] Parity is also violated individually, but since the Standard Model is a chiral theory anyway, P violation is a less remarkable feature than CP violation.

[4] The reason why QCD also conserves CP if $\theta = \pi$ is that CP maps θ from π to $-\pi$, which is precisely a shift by the periodicity 2π of the θ-angle [36].

[5] Since CP is also conserved at $\theta = \pi$, one may think that θ could also be very close to π. However, as Vafa and Witten argued in ref. [46], this cannot be the case due to the sign of quark condensates.

While being an impressive experimental achievement, this constitutes a serious problem of the SM. All dimensionless parameters like θ are naturally expected to be of order unity and none of the others have been proven to be this small. A good theory should provide an intrinsic reason why θ is precisely set to a specific value. Such an unexpected coincidence is called a fine-tuning problem. Furthermore, the smallness of θ cannot be called natural by 't Hooft's definition [47] because setting θ to zero does not enhance the symmetry of the theory. The reason for this is that CP is already independently violated in the weak sector of the SM by a non-vanishing phase of the CKM matrix [48]. In conclusion, the SM does not explain why the strong interaction is CP-symmetric to such a high precision. This constitutes the strong CP problem of the SM.

The problem becomes even more puzzling when the full particle content of the SM is considered. Besides the vacuum angle of QCD, the theta-term receives another contribution from the phases of quark masses due to the chiral anomaly. If a gauge theory includes at least one chiral fermion, it can be anomalous under a global chiral rotation of these fermionic fields. A toy example would be the theory defined by the Lagrangian

$$\mathcal{L} = -\frac{1}{2}\mathrm{Tr}\, G_{\mu\nu}G^{\mu\nu} + \bar{\psi}(i\not{D} + m)\psi \; . \tag{2.16}$$

Under a chiral transformation, the mass term receives a phase. Furthermore, the fermionic path integral is not invariant, which adds a topological term to the Lagrangian (see e.g. ref. [29]).

$$\psi_L \to e^{i\alpha/2}\psi_L \quad , \quad \psi_R \to e^{-i\alpha/2}\psi_R \tag{2.17}$$

$$\Rightarrow \quad \mathcal{L} \to -\frac{1}{2}\mathrm{Tr}\, G_{\mu\nu}G^{\mu\nu} + \bar{\psi}(i\not{D} + me^{-i\alpha})\psi - \alpha\frac{g_s^2}{16\pi^2}\mathrm{Tr}\, G_{\mu\nu}\tilde{G}^{\mu\nu}. \tag{2.18}$$

This is the famous Adler–Bell–Jackiw anomaly [49, 50]. It also appears in the SM, where all quarks gain masses through the Higgs mechanism and it is common to transform the fields into a mass eigenbasis.[6] To do this, the Yukawa couplings are diagonalized by unitary vector transformations. If the resulting diagonal mass matrix M_q has a common phase, axial rotations are required to end up with a real diagonal quark mass matrix. This means that the actual physically observable angle in the SM, here denoted $\bar{\theta}$, is the sum of θ and the argument of the determinant of the quark mass matrix[7],

$$\bar{\theta} = \theta + \mathrm{Arg}(\mathrm{Det}(M_q)). \tag{2.19}$$

To explain the smallness of the nEDM in the SM, $\bar{\theta}$ must be smaller than 1.5×10^{-10}. This requires an almost exact cancellation between two independent and a priori unrelated parameters, which seems unlikely to occur by chance and begs for a theoretical explanation.

---

[6] Note that if the theory contains a massless coloured fermion, the chiral rotation would be a symmetry of the Lagrangian and the theta-angle could be rotated away. This is why a massless up quark was considered a possible solution to the strong CP problem but recent lattice calculations point towards a finite mass [51–54].

[7] In principle, such rotations also give rise to theta-terms of the other SM gauge groups. However, the fermion content and particularly the chiral nature of the electroweak interaction allows to rotate the electroweak theta angles to zero in the SM [5, 55–57].

2.3 Peccei–Quinn solution

A dynamical solution to the strong CP problem was proposed by Roberto Peccei and Helen Quinn in 1977 [14, 58]. Their idea is based on the fact that in the presence of instantons[8] the vacuum energy depends on $\bar{\theta}$. When $\bar{\theta}$ is regarded as a dynamical parameter, this corresponds to a potential $V(\bar{\theta})$. In addition, they found that this potential has two stationary points at $\bar{\theta} = 0$ and π, the former of which is a true minimum. Hence, observing CP conservation in strong interactions merely corresponds to finding $\bar{\theta}$ at or close to the minimum of its potential, which is a natural expectation. In this sense, Peccei and Quinn discovered that the strong CP problem solves itself if it is allowed to. The only requirement is a dynamical $\bar{\theta}$, which can relax to its minimum.

Shortly after this discovery by Peccei and Quinn, Frank Wilczek and Steven Weinberg noted independently of one another that this solution to the strong CP problem indicates the existence of a new particle [12, 13] because the dynamical angle corresponds to a new scalar field. Wilczek called this new particle axion [59]. The name stems from a laundry detergent, which is fitting because the axion is devised to wash the SM model clean of one of its stains and because it is related to the anomalous breaking of an *axial* symmetry.

On the formal level, the solution described above amounts to introducing a new scalar field a, which couples to the $G\widetilde{G}$-term in exactly the same way as $\bar{\theta}$,

$$\mathcal{L}_{\bar{\theta},a} = \frac{1}{2}\left(\partial_\mu a\right)\left(\partial^\mu a\right) + \frac{g_s^2}{16\pi^2}\left(\bar{\theta} + \frac{a}{f_a}\right)\operatorname{Tr} G_{\mu\nu}\widetilde{G}^{\mu\nu}\ . \tag{2.20}$$

How these terms can be generated from spontaneous breaking of a global U(1) symmetry is discussed in chapter 3. The scalar field is written with a canonical kinetic term. Therefore, the interaction with $G\widetilde{G}$ must come with a new scale called the axion decay constant f_a. The fraction a/f_a takes the role of the dynamical theta-angle. The fact that the global minimum of the effective potential for a must be at the CP conserving value, i.e. at $a = -\bar{\theta} f_a$, can be most generically seen following an argument by Vafa and Witten [46]. Based on the assumption that the only imaginary term in the Euclidean path integral comes from $G\widetilde{G}$, the following inequality for the effective potential has to apply,

$$\begin{aligned}
&\exp\left(-V_{\text{eff}}\left[\bar{\theta} + \frac{a}{f_a}\right]\right) \\
&= \left|\int \mathcal{D}G^\mu \exp\left(-S_{\text{eff}}[\phi, G^\mu] - i\int \mathrm{d}^4x_{\mathrm{E}} \frac{g_s^2}{16\pi^2}\left(\bar{\theta} + \frac{a}{f_a}\right)\operatorname{Tr} G_{\mu\nu}\widetilde{G}^{\mu\nu}\right)\right| && (2.21) \\
&\leq \int \mathcal{D}G^\mu \left|\exp\left(-S_{\text{eff}}[\phi, G^\mu] - i\int \mathrm{d}^4x_{\mathrm{E}} \frac{g_s^2}{16\pi^2}\left(\bar{\theta} + \frac{a}{f_a}\right)\operatorname{Tr} G_{\mu\nu}\widetilde{G}^{\mu\nu}\right)\right| && (2.22) \\
&= \int \mathcal{D}G^\mu \left|\exp\left(-S_{\text{eff}}[\phi, G^\mu]\right)\right| && (2.23) \\
&= \exp\left(-V_{\text{eff}}[0]\right) && (2.24)
\end{aligned}$$

---

[8] Originally, the first paper on the PQ solution was titled 'CP conservation in the presence of instantons'. The journal insisted to replace 'instantons' by 'pseudoparticles' even though instantons should better not be considered as particles for reasons discussed in section 2.1.

$$\Rightarrow \quad V_{\text{eff}}[0] \leq V_{\text{eff}}\left[\bar{\theta} + \frac{a}{f_a}\right] . \tag{2.25}$$

This means that the minimum of the potential will be at the CP conserving value where $\bar{\theta} + a/f_a = 0$ and no spontaneous breaking of CP can occur. However, the assumption made above is not strictly fulfilled in the SM. One can show that CP violating terms will necessarily be imaginary in the Euclidean action [60]. The CP violating phase of the CKM matrix is such a term and it will indeed contribute to the effective axion potential by shifting its minimum away from the CP conserving value. In order not to spoil the PQ solution, this contribution has to be tiny. In fact this is the case since the CKM phase is only a physical parameter if the SM includes at least three generations of quarks [61]. Hence, all three generations must appear in Feynman diagrams[9] needed for the calculation of the effective axion potential or the resulting nEDM from the CKM phase. These kinds of diagrams are highly suppressed by at least two loops and a heavy quark mass. This is why the nEDM in the SM is several orders of magnitude smaller than the current limit in equation (2.13) and why the CKM phase does not spoil the argument above [62–64]. For similar reasons, a non-vanishing phase of the PMNS matrix would not be an issue. Secondly, there is P violation in the SM because only left-handed particles interact weakly. This also contributes to the imaginary part of the Euclidean action, but is no source of CP violation and the position of the vacuum is not shifted [14, 58].

Apart from the position of its minimum, the specific form of the axion potential is of great interest for instance for the DM phenomenology of axions. Most importantly, the second derivative at the minimum of the potential defines the axion mass as [54]

$$m_a^2 = \frac{\delta^2}{\delta a^2} \log \mathcal{Z}\left(\bar{\theta} + \frac{a}{f_a}\right)\bigg|_{a=\langle a\rangle} = \frac{1}{f_a^2}\frac{d^2}{d\theta^2}\log \mathcal{Z}(\theta)\bigg|_{\theta=0} = \frac{\chi_{\text{top}}}{f_a^2} . \tag{2.26}$$

$\mathcal{Z}$ denotes the generating functional, which depends on a dynamical axion field and a constant theta-term. In the last step, the topological susceptibility χ_{top} was defined. The QCD axion mass is thus proportional to f_a^{-1}. To leading order in chiral perturbation theory, the axion mass can be expressed in terms of QCD parameters [13, 54] as[10]

$$m_a^2 \simeq \frac{m_u m_d}{(m_u + m_d)^2}\frac{m_\pi^2 f_\pi^2}{f_a^2} . \tag{2.27}$$

m_u and m_d are the masses of up and down quarks. m_π is the pion mass and f_π its decay constant. This result reflects the fact that θ is unphysical when one of the quarks is massless. In this case, the axion field does not obtain a mass and CP is conserved for any value of a. More recently, the most accurate calculations of χ_{top} and m_a were performed by combining next-to-next-to-leading order calculations in

---

[9] Unlike the potential and nEDM generated by the theta-term, the contributions from the phase of the CKM matrix can be computed perturbatively.

[10] See refs. [65–68] for some recent work motivating the possibility that QCD axions could significantly deviate from this relation and be heavier than expected.

chiral perturbation theory [54, 69] with lattice evaluations of QCD parameters like the mass ratios of light quarks [70–73]. The numerical result is [69]

$$m_a = \frac{\sqrt{\chi_{\text{top}}}}{f_a} = 5.69(5)\,\mu\text{eV}\left(\frac{10^{12}\,\text{GeV}}{f_a}\right) . \tag{2.28}$$

Evaluations of the full axion potential are even more involved. Again, only the leading-order result is cited here [54],

$$V(a) \simeq -m_\pi^2 f_\pi^2 \sqrt{1 - \frac{4m_u m_d}{(m_u + m_d)^2}\sin^2(\frac{a}{2f_a})} , \tag{2.29}$$

where the axion field was shifted such that the minimum lies at $a = 0$. More accurate evaluations of the potential and finite temperature effects, which are relevant in the early universe, can be found in ref. [54].

In summary, the Peccei–Quinn solution of the strong CP problem proposes to add a dynamical scalar field, which couples to $G\widetilde{G}$ and thus takes the role of a dynamical theta-angle. This new scalar is called the *QCD axion*. Non-perturbative QCD effects induce a potential, whose minimum lies at the CP conserving value. Potential and mass of the axion only depend on a single free parameter – the axion decay constant f_a, which acts as a suppression scale of all axion interaction terms.

2.4 Axion dark matter

The sole original motivation for the axion was to solve the strong CP problem [12, 13]. Nevertheless, shortly after its proposition, it was realised that the axion could be a natural DM candidate [15, 16]. By construction, the axion is weakly coupled to SM particles, light and therefore long-lived. These are ideal properties of a DM particle. Before giving an overview of the dynamics of the cosmological axion field, the main arguments for the existence of dark matter and its nature are recalled in the following. Extensive books and reviews on this topic can be found for example in refs. [74–77].

2.4.1 The dark matter puzzle

DM is a conjectured energy abundance in the universe which at most reacts very weakly with the ordinary matter included in the SM. It is only known to interact gravitationally and its energy density should scale with the inverse volume, i.e. like ordinary non-relativistic matter. The observational evidence for its existence is overwhelming and consistently points towards a cosmological abundance which is more than five times as large as the one of ordinary baryonic matter [75]. Cosmological abundances are typically quantified by the density parameter Ω_i, which is the density of some cosmological species i relative to the critical density of the universe. The Planck collaboration found [78]

$$\Omega_{\text{b}} = 0.049 , \quad \Omega_{\text{DM}} = 0.259 , \quad \Omega_\Lambda = 0.691 , \tag{2.30}$$

where b, DM and Λ label baryonic matter, dark matter and dark energy, respectively.[11] These numbers lead to the often-cited conclusion that modern physics only understands and is able to describe the nature of about 5 % of the total energy content of the universe.

One of the oldest observations that lead to the proposition of DM are galactic rotation curves [79]. In the outer regions of galaxies, stars are observed to orbit at a velocity much larger than what would be expected from the enclosed mass inside its orbit. This can be explained by a larger gravitational pull coming from a different type of matter which escapes observation but nevertheless contributes to the mass of the galaxy [75].

Another powerful technique for observing matter distributions purely by its gravitational impact are called gravitational lenses (see ref. [80] for a review on this topic). This subsumes a number of effects which can occur when light coming from a distant object encounters strong gravitational potentials. Depending on the exact matter distribution, this can lead to image distortions like shear, multiple images of the same objects, Einstein rings or an increase in brightness [75]. From these effects, the matter distribution between the object and the observer can be reconstructed. Consistently, such observations point towards galaxies being more massive compared to the expectation based on the amount of observable baryonic matter.

A singular piece of compelling evidence for DM is the famous observation of the Bullet Cluster [81]. It formed from the collision of two smaller galaxy clusters. By observations in the x-ray as well as in the visible spectrum, both the baryonic and the overall matter density can be mapped. The latter is only detectable through its gravitational interaction, i.e. via its lensing effects. Intriguingly, one can observe how the centre of mass of the two smaller clusters passed through one another without significant interference thus proving a lack of interaction of the majority of the matter involved. The baryonic matter, on the other hand, slowed down due to interactions with the collision partner. All of these observations are in perfect agreement with a cluster consisting mainly out of collisionless dark matter. The Bullet Cluster is particularly hard to explain in modified gravity theories since a clear differentiation between different kinds of matter abundances can be observed [75, 81].

Finally, the inclusion of DM is a crucial ingredient in our current understanding of the history of the universe. This fact becomes most apparent in observations of the cosmic microwave background (CMB). This earliest detectable light exhibits the structures in the universe at last scattering and it is a powerful probe of cosmological parameters. The interplay between these early structures and the resulting CMB power spectrum is far from trivial. An overview of the important effects can for instance be found in ref. [74]. Overall, these observations are in extremely good agreement[12] with the ΛCDM model of the universe when the density parameters are set to the values stated in equation (2.30) [78].

In conclusion, several independent observations point towards a large, unknown

---

[11] Only rounded central values are quoted here because they are sufficient for the point of the argument.

[12] Potential deviations from the ΛCDM model like the cusp-core problem [82, 83], the missing satellite problem [84, 85] or the too-big-to-fail problem [86] have mostly been successfully solved by the inclusion of baryonic effects in the formation of galaxies [87–89].

and weakly interacting population of matter being present in our universe. The question, which is one of the most pressing puzzles in physics today, is what this matter consists of. The only objects known to exist and which have the right properties are SM neutrinos and black holes. The former cannot comprise the whole amount of DM due to neutrinos being very light and fermionic. The latter remains a valid option but only for relatively broad distributions of masses [90, 91]. This is why many candidates beyond the SM are also considered as DM. The key properties of a good candidate are weak couplings to the SM and the existence of a valid production mechanism in the early universe leading to the right abundance today.

2.4.2 Production of axion dark matter

The axion is a typical example for a light DM candidate, which means that it has to be bosonic. Otherwise Pauli pressure would prohibit the existence of a large and cool enough abundance [92]. Its lightness also means that it cannot be produced thermally. If the axions made up DM and had been in thermal equilibrium in the early universe, structure formation to the extent observed today would have not been possible [93]. Fortunately, the axion generically provides a non-thermal production mechanism, which can lead to the right abundance today – the misalignment mechanism. Even though several versions with slightly different properties exist, the basic idea is always the same.

The description of the axion as in equation (2.20) becomes valid once a new global symmetry – the PQ symmetry – is spontaneously broken. This is explained in detail in chapter 3. The evolution of the classical axion field is then governed by the Klein–Gordon-equation in an FLRW metric [94],

$$\ddot{a} + 3H\dot{a} + V'(a) = 0 \ , \tag{2.31}$$

where a denotes the axion field[13], H is the Hubble parameter and V the periodic axion potential. As long as H is much larger than the potential term, friction dominates and the field is effectively constant. But once H drops to values much smaller than the axion mass, $V'(a) \simeq m_a^2 a$, the field starts to oscillate. These spatially coherent oscillations have an energy density which scales like matter with the cubic inverse of the cosmological scale factor. Because of spatial homogeneity (at least of the classical background field) the velocity dispersion is almost zero, the produced abundance of axions is extremely cold and is clearly not in thermal equilibrium with the dominating radiation at the time.

The initial conditions of this process can be set in either of two ways depending on whether the PQ symmetry is broken during or after inflation. During SSB, the axion field takes a random value at every point in space, because there is no preferred position on the S^1 vacuum manifold. If PQ breaking occurs during inflation, the random field is stretched out such that the axion field has a constant value in the whole observable universe. While this is a very simple scenario, the value of the

---

[13] In the cosmological context the axion field is often denoted A or ϕ in order to avoid confusion with the cosmological scale factor a. This is not necessary here since the cosmological scale factor only appears implicitly through H.

constant initial value of $\theta_i = a_{\text{init}}/f_a$ in the pre-inflationary scenario is unknown because it can randomly take any value between 0 and 2π. The resulting dark matter abundance depends on the subsequent expansion history of the universe and the axion decay constant f_a. To give a rough estimate, today's abundance can be approximated by [94]

$$\Omega_a \sim 4 \times 10^4 \left(\frac{f_a}{10^{16}\,\text{GeV}} \right)^{\frac{7}{6}} \theta_i^2 \,, \tag{2.32}$$

in the case of $f_a \lesssim 2 \times 10^{15}\,\text{GeV}$. This condition ensures that the oscillations in the axion field start while $T > \Lambda_{\text{QCD}}$ and thus the axion mass can be computed in the dilute instanton gas approximation. For $f_a \gtrsim 2 \times 10^{17}\,\text{GeV}$, the relation becomes [94]

$$\Omega_a \sim 1 \times 10^4 \left(\frac{f_a}{10^{16}\,\text{GeV}} \right)^{\frac{3}{2}} \theta_i^2 \,, \tag{2.33}$$

where the axion mass has reached its zero temperature value at the onset of oscillations.[14] Both of these results were derived in the harmonic approximation of the axion potential, $V(a) = \frac{1}{2} m_a^2 a^2$, and the results receive significant corrections for large initial field values. It becomes clear from the equations that a lighter axion, i.e. with larger f_a, is produced in larger abundances. This can eventually lead to overproduction and therefore an overclosed universe. A smaller abundance can always be explained by a small initial misalignment angle θ_i but only at the cost of fine-tuning.

The alternative, the post-inflationary scenario, is in general more predictive. When the PQ symmetry is broken after inflation, the axion field takes a random value in each causally disconnected patch of the observable universe. While still random, the probability distribution is predicted by the theory and can be connected to a prediction of the total amount of produced dark matter. This means that the misalignment angle θ_i^2 in equations (2.32) and (2.33) can simply be replaced by its expectation value $\langle \theta_i^2 \rangle$. However, there is a significant complication in the post-inflationary scenario. The initial values can make a whole winding around the S^1 field manifold along a closed path in space. This results in topological defects called axion strings. These strings carry a significant amount of the field's energy and eventually decay while the universe expands. The calculation of the resulting axion DM abundance is numerically challenging. State-of-the-art calculations [96–100] still differ significantly but point towards an axion mass of tens or hundreds of µeV if the axion is supposed to make up the entire DM in the universe.

In conclusion, a realistic estimate of the viable mass range of QCD axion DM is $1\,\mu\text{eV} \lesssim m_a \lesssim 200\,\mu\text{eV}$. As mentioned above, a lighter mass would always be possible in the pre-inflationary scenario if some tuning of the initial misalignment angle is accepted. This is often referred to as the *anthropic window*. The upper bound comes from the larger abundance estimates in the post-inflationary scenario. A more solid

---

[14] In the intermediate regime, $2 \times 10^{15}\,\text{GeV} \lesssim f_a \lesssim 2 \times 10^{17}\,\text{GeV}$, the oscillations start during or close to the QCD phase transition and neither of the two expressions above is accurate [95].

constraint in the same mass range for QCD axions can be derived from non-thermal production via the axion-gluon interaction (see ref. [101] for the latest evaluation of such a bound). It should also be mentioned that with some modifications, the standard mass-abundance relations can be modified by several orders of magnitude. Recent proposals for this kind of mechanism include a dynamical PQ scale [102] or thermal friction from additional gauge fields [103, 104]. So even though there is a preferred mass range, where the axion DM production is realised in its simplest form, one should try to search in an even wider range in order to rule out non-standard scenarios.

3 Axion models and effective field theories

The axion, as it was introduced in the previous sections, is essentially a dynamical theta-angle of QCD. Its coupling of the form $a\,G\widetilde{G}$ together with its shift symmetry, are all that is needed to solve the strong CP problem. However, this coupling term has mass dimension five, which makes it in general non-renormalizable. Consequently, it has to be an effective operator of a more fundamental UV-complete theory. These UV completions can take various forms and are referred to as axion models. Even though the underlying theories vary, the solution of the strong CP problem always follows the same logic.

The important property of shift invariance can be realised by introducing the axion as a Goldstone boson of a spontaneously broken global symmetry, the $\mathrm{U}(1)_{\mathrm{PQ}}$. Before looking at specific examples of how such a new symmetry can be realised, it is instructive to systematically think about the minimal required field content of such a theory. For the generation of the axion-gluon coupling, $\mathrm{U}(1)_{\mathrm{PQ}}$ must be a chiral symmetry under which at least one coloured fermion is charged. Such a symmetry does not exist in the SM since the only global symmetry (at the Lagrangian level) involving quarks is the one related to baryon number conservation, which is not chiral. It was therefore argued in ref. [105] that any axion model will need to introduce either additional fermions or scalars in order to extend the space of possible charges under $\mathrm{U}(1)_{\mathrm{PQ}}$. Thus, the axion will never be the only additional particle beyond the SM in a UV-complete theory.

Throughout this thesis, charges under $\mathrm{U}(1)_{\mathrm{PQ}}$ of a particle p (with fixed chirality) will be denoted χ_p. A fermionic field $\psi_{L/R}$ and a scalar field ϕ would therefore transform as

$$\psi_L \to e^{i\alpha\chi_{\psi_L}}\psi_L \quad , \quad \psi_R \to e^{i\alpha\chi_{\psi_R}}\psi_R \quad , \quad \phi \to e^{i\alpha\chi_\phi}\phi \; . \tag{3.1}$$

Exactly as described in section 2.2, this symmetry can feature an anomaly which will result in additional terms in the Lagrangian. For fermions charged under the $\mathrm{SU}(3)_\mathrm{c}$ and $\mathrm{U}(1)_{\mathrm{EM}}$ gauge symmetries of the SM[15], these are [14, 49, 50, 58, 107]

$$\mathcal{L}_{\text{anomaly}} = -\alpha N \frac{g_s^2}{32\pi^2} G^a_{\mu\nu}\widetilde{G}^{a\mu\nu} - \alpha E \frac{g_s^2}{32\pi^2} F_{\mu\nu}\widetilde{F}^{\mu\nu} \; , \tag{3.2}$$

where N and E are the anomaly coefficients defined as

$$N = \sum_f (\chi_{f_L} - \chi_{f_R}) T(R_f) \tag{3.3}$$

$$E = 2\sum_f (\chi_{f_L} - \chi_{f_R}) Q_f^2 \; . \tag{3.4}$$

---

[15] Couplings to weak gauge fields are also generated but are of less phenomenological interest and are not further investigated in this thesis. See ref. [106] for a detailed discussion of these terms.

$T(R_f)$ denotes the Dynkin index of the $\mathrm{SU}(3)_\mathrm{c}$ representation R_f, which is unity in the case of a fundamental representation[16], and Q_f are the electric charges. The sums run over all irreducible representations of the gauge group in question. In the axion models presented below, it will be necessary to perform chiral rotations which depend on the axion field. In these cases, the same transformation rules apply and the rotation angle α is simply replaced by a/f_a. The popular axion models discussed in this chapter will later serve as benchmarks for phenomenological studies. This is followed by an introduction into axion effective field theories (EFTs) which provide a model independent parametrisation of all possible low-energy effective interactions.

3.1 KSVZ-type models

Hadronic axion models are the subset of models which extend the vector space of PQ charges by introducing new fermionic fields. The simplest version of this is the prominent KSVZ model [109, 110], which contains the minimal field content required for an anomalous breaking of the $\mathrm{U}(1)_\mathrm{PQ}$ symmetry. This includes a fermion Q charged under both the chiral $\mathrm{U}(1)_\mathrm{PQ}$ symmetry and the $\mathrm{SU}(3)_\mathrm{c}$ gauge symmetry, which couples to the PQ scalar Φ via a standard Yukawa interaction. The full KSVZ Lagrangian is therefore given by

$$\mathcal{L}_\mathrm{KSVZ} = i\bar{Q}\not{D}Q + (\partial_\mu\Phi)(\partial^\mu\Phi) - V(|\Phi|^2) - (y\Phi\bar{Q}_L Q_R + \text{h.c.}) \ . \tag{3.5}$$

The model then assumes that the complex scalar acquires a vacuum expectation value (VEV) v, spontaneously breaking $\mathrm{U}(1)_\mathrm{PQ}$. Inserting the convenient parametrisation

$$\Phi = \frac{1}{\sqrt{2}}\,(v+\rho)\,e^{i\frac{a}{v}}, \tag{3.6}$$

the resulting mass of the quark is given by $m_Q = yv/\sqrt{2}$. The radial mode ρ acquires a mass of the same scale, while the Goldstone mode a remains massless. The desired coupling of a to gluons is found most easily by performing a chiral field redefinition,

$$Q \to e^{i\frac{a}{2v}\gamma^5} Q \ . \tag{3.7}$$

This eliminates the dependence on a from the Yukawa terms but instead induces a QCD axion coupling,

$$\mathcal{L} \supset -\frac{\alpha_s}{8\pi}\frac{a}{v} G^a_{\mu\nu}\widetilde{G}^{a\mu\nu} = -\frac{\alpha_s}{8\pi}\frac{a}{f_a} G^a_{\mu\nu}\widetilde{G}^{a\mu\nu} \ , \tag{3.8}$$

where v was identified as the axion decay constant f_a in this model. After integrating out the heavy fields Q and ρ, this is the only remaining axion coupling to the SM. This pure KSVZ model is therefore not only minimal in its field content

---

[16] Note that some authors like e.g. [108] use a different convention where $T(R_f) = 1/2$ for a fundamental representation. They also omit the factor of 2 in the definition of E and thus arrive at the same value for E/N. For the purpose of chapter 5, a convention in which N is always an integer is chosen.

but also results in the minimal axion couplings required for solving the strong CP problem. However, this simplicity comes at a price. Because the heavy quark Q only interacts via the strong force, it is stable on cosmological time scales [109] and the population of thermal relics from the Big Bang would be observable today [111–113]. Solving this problem leads to the wider class of KSVZ-type or hadronic axion models. One can allow Q to have other gauge charges, e.g. a non-vanishing hypercharge, which means that it can mix with SM quarks and will decay. Such a model would contain additional axion couplings to SM gauge bosons because the chiral rotation generates anomalous coupling terms for all gauge groups that the fermion is charged under. In general, hadronic axion models can even feature a number of these coloured fermions in various representations of the gauge groups. The resulting axion interactions with gluons and photons can then be derived from the anomaly coefficients in equations (3.3) and (3.4).[17] A comprehensive list of suitable representations in hadronic axion models with a single heavy quark is described in refs. [108, 115].

3.2 DFSZ-type models

An alternative to introducing fermionic fields charged under the new $\mathrm{U}(1)_{\mathrm{PQ}}$ symmetry is the expansion of the scalar sector. More specifically, one can consider a two Higgs doublet model (2HDM) and identify the axion as a linear combination of the two pseudoscalar Goldstone modes. This was exactly the idea of the original PQWW axion model [12, 13]. In this scenario the axion decay constant is closely related to the electroweak scale. The model was quickly ruled out by experimental observations because it predicted axion couplings of the same order as the weak interactions. However, the 2HDM approach can be tweaked to avoid these constraints by introducing another complex scalar. This is the well-known DFSZ model [116, 117], which will be introduced in more detail below. The description roughly follows the one in ref. [118].

In a two Higgs doublet model, the up- and down-type quarks couple to different Higgs doublets. The Yukawa terms can therefore be written as

$$\mathcal{L}_Y = -\lambda_u \bar{q}_L H_u u_R - \lambda_d \bar{q}_L H_d d_R - \lambda_e \bar{L} H_d e_R + \text{h.c.}\ . \tag{3.9}$$

Exactly as in the Standard Model, the fermions obtain their masses through the VEVs of the neutral components of the Higgs doublets. In the lepton Yukawa term, H_d could be replaced by ϵH_u^*. These two possibilities of generating lepton masses through H_d or H_u are called DFSZ I and II, respectively. For simplicity, only the former is presented here, but all steps are completely equivalent in both cases. The PQ charge of the two Higgs doublets are identical, but they differ in their hypercharge. This means that the quark Yukawa terms shown above are the only ones compatible with all symmetries.

---

[17] Directly using the anomaly coefficients as coupling constants only works for non-chiral gauge groups like $\mathrm{SU}(3)_\mathrm{c}$ or $\mathrm{U}(1)_{\mathrm{EM}}$ and only in the limit of infinite m_Q. This point is discussed in detail in ref. [114].

The only other BSM field in the DFSZ model is a complex scalar Φ, which couples to the doublets via the term $H_d^\dagger \epsilon H_u^* \Phi^n + \text{h.c.}$. Renormalizability requires n to be either 2 or 1. When all three scalar fields acquire a VEV, the neutral Higgs components $H^0_{u/d}$ and Φ can be expressed as

$$H^0_{u/d} = \frac{1}{\sqrt{2}} \left(v_{u/d} + \rho_{u/d} \right) e^{i\theta_{u/d}} \quad \text{and} \quad \Phi = \frac{1}{\sqrt{2}} \left(v_\Phi + \rho_\Phi \right) e^{i\theta_\Phi} \,. \tag{3.10}$$

Here, v, ρ and θ denote the VEVs, radial and angular modes, respectively. In this parametrisation, the three possible neutral Goldstone fields $a_{u/d/\Phi} = v_{u/d/\Phi}\theta_{u/d/\Phi}$ all obtain a canonical kinetic term. In the following, only orthogonal rotations in this field basis are made in order to keep the kinetic terms diagonal.

The subspace spanned by a_d and a_u is the one affected by hypercharge gauge transformations, but the combination $v_d a_u + v_u a_d$ is invariant under $\mathrm{U}(1)_\mathrm{Y}$. This motivates the definitions

$$\begin{pmatrix} a_w \\ a_Z \end{pmatrix} = \begin{pmatrix} \cos\beta & \sin\beta \\ -\sin\beta & \cos\beta \end{pmatrix} \begin{pmatrix} a_u \\ a_d \end{pmatrix} \quad \text{with} \quad \tan\beta = \frac{v_u}{v_d}, \tag{3.11}$$

where a_w is the PQWW axion.[18] This direction is invariant under gauge transformations while the Goldstone a_Z is the one which becomes the longitudinal mode of the Z boson in unitary gauge. It is convenient to define $v_F^2 = v_u^2 + v_d^2$ such that $\cos\beta = v_d/v_F$ and $\sin\beta = v_u/v_F$.

The coupling between the Higgs doublets and Φ generates a non-vanishing potential for the pseudoscalar modes. When all three scalars acquire a VEV and the heavy radial modes are neglected, the potential becomes

$$\begin{aligned} &\lambda H_d^\dagger \epsilon H_u^* \Phi^n + \text{h.c.} \\ \xrightarrow{\text{SSB}} \quad & |\lambda| \frac{v_u v_d v_\Phi^n}{2^{\frac{2+n}{2}}} e^{i(n\theta_\Phi - \theta_u - \theta_d + \delta)} + \text{h.c.} = |\lambda| \frac{v_u v_d v_\Phi^n}{2^{\frac{n}{2}}} \cos\left(\frac{n}{v_\Phi} a_\Phi - \frac{v_F}{v_u v_d} a_w + \delta\right) . \end{aligned} \tag{3.12}$$

λ is an arbitrary complex coupling constant with argument δ. The potential is independent of a_Z, which is an immediate consequence of the gauge invariance of the interaction term. One can now rotate into a mass diagonal basis,

$$\begin{pmatrix} a \\ a_m \end{pmatrix} = \begin{pmatrix} \cos\gamma & \sin\gamma \\ -\sin\gamma & \cos\gamma \end{pmatrix} \begin{pmatrix} a_\Phi \\ a_w \end{pmatrix} \quad \text{with} \quad \tan\gamma = \frac{n v_u v_d}{v_F v_\Phi} = \frac{n v_w}{v_\Phi} \,, \tag{3.13}$$

where the scale $v_w = v_u v_d / v_F$ was introduced. In this basis, the argument of the cosine in the scalar potential in equation (3.12) only depends on the massive direction a_m, while the DFSZ axion a is the only pseudoscalar direction which remains massless and is not eaten by a gauge boson.

The couplings of the DFSZ axion to gauge bosons are finally obtained by performing a chiral rotation of fermions, eliminating the phase of the Yukawa terms in

---

[18] In a scenario without the additional complex scalar Φ, a_w remains massless up to the anomalous symmetry breaking.

equation (3.9),

$$\begin{pmatrix} u_L \\ d_L \\ e_L \end{pmatrix} \to \begin{pmatrix} e^{i\frac{\theta_u}{2}} u_L \\ e^{i\frac{\theta_d}{2}} d_L \\ e^{i\frac{\theta_d}{2}} e_L \end{pmatrix} \qquad \begin{pmatrix} u_R \\ d_R \\ e_R \end{pmatrix} \to \begin{pmatrix} e^{-i\frac{\theta_u}{2}} u_R \\ e^{-i\frac{\theta_d}{2}} d_R \\ e^{-i\frac{\theta_d}{2}} e_R \end{pmatrix} . \tag{3.14}$$

Assuming that a_Z is set to zero by a gauge transformation and that that the heavy field a_m is integrated out, the phases $\theta_{u/d}$ only depend on the axion field a,

$$\theta_u = \frac{\cos^2\!\beta}{v_w} a_w = \frac{\cos^2\!\beta \, \sin\gamma}{v_w} a = \frac{\cos^2\!\beta}{3} \frac{3n}{\sqrt{n v_w^2 + v_\Phi^2}} a = \frac{\cos^2\!\beta}{3 f_a} a \, , \tag{3.15}$$

where f_a was defined as

$$f_a = \frac{\sqrt{(n v_w)^2 + v_\Phi^2}}{3n} \, . \tag{3.16}$$

Analogously, one finds $\theta_d = a \sin^2(\beta)/(3 f_a)$. The chiral rotation will eliminate the axion from the Yukawa terms and instead induce derivative couplings to fermions. Additionally, the chiral anomaly generates couplings to gauge bosons as described in equations (3.3) and (3.4). Accounting for all generations of fermions, the resulting terms in the Lagrangian are[19]

$$\begin{aligned} \mathcal{L}_{\mathrm{DFSZ}} = & - \frac{\alpha_s}{8\pi} G^a_{\mu\nu} \widetilde{G}^{a\mu\nu} \, 3(\theta_u + \theta_d) \\ & - \frac{\alpha_{\mathrm{EM}}}{8\pi} F_{\mu\nu} \widetilde{F}^{\mu\nu} \, 2 \cdot 3 \left[3 \left(\frac{2}{3} \right)^2 \theta_u + 3 \left(\frac{-1}{3} \right)^2 \theta_d + (-1)^2 \theta_d \right] \\ & + \bar{u} \gamma^\mu \gamma^5 u \frac{\partial_\mu \theta_u}{2} + \left(\bar{d} \gamma^\mu \gamma^5 d + \bar{e} \gamma^\mu \gamma^5 e \right) \frac{\partial_\mu \theta_d}{2} \end{aligned} \tag{3.17}$$

$$\begin{aligned} = & - \frac{\alpha_s}{8\pi} \frac{a}{f_a} G^a_{\mu\nu} \widetilde{G}^{a\mu\nu} - \frac{\alpha_{\mathrm{EM}}}{8\pi} \frac{8}{3} \frac{a}{f_a} F_{\mu\nu} \widetilde{F}^{\mu\nu} \\ & + \frac{\sin^2\!\beta}{3} \bar{u} \gamma^\mu \gamma^5 u \frac{\partial_\mu a}{2 f_a} + \frac{\cos^2\!\beta}{3} \left(\bar{d} \gamma^\mu \gamma^5 d + \bar{e} \gamma^\mu \gamma^5 e \right) \frac{\partial_\mu a}{2 f_a} \, . \end{aligned} \tag{3.18}$$

At this point, the reason for the normalization of f_a that was chosen above becomes clear since a/f_a takes the role of the dynamical theta-angle and the induced axion potential has a discrete shift symmetry $a \to a + 2\pi f_a$. Accordingly, the axion-photon coupling is given by $E/N = 8/3$. The choice of n only affects the dependence of f_a on the VEVs of the three scalar fields. Doing the same steps as above for the DFSZ II model, one arrives at the same Lagrangian but with factors of $-\sin^2(\beta)/3$ for the lepton coupling and $E/N = 2/3$ for the photon term.

The DFSZ models type I and II often serve as benchmark axion models, especially when couplings to SM fermions are required. Additionally, there is a large class of similar models that also employ additional Higgs doublets but lead to different

---

[19] Note that in a quark mass eigenbasis, the different charges of up- and down-type quarks also leads to terms of the type $e^{ia/f_a} W_\mu \bar{u}_L \gamma^\mu d_L$. The corresponding four particle vertex is however suppressed by both the weak coupling and f_a^{-1} and it is therefore commonly neglected.

effective axion couplings. The general logic is always the same, but for instance additional Higgs doublets are included [119, 120] or PQ charges differ between the three fermion generations [119, 121, 122]. These models are often referred to as DFSZ-type. At the expense of some fine-tuning, they can be designed to avoid astrophysical constraints [119, 121, 123]. Most importantly, their existence motivates to search for axions via all possible couplings and to not merely constrain the one low-dimensional parameter space of a specific model.

3.3 Non-canonical axion models

Besides the two well-known models and their modifications, which were presented above, there is a vast landscape of conceivable QCD axion models [124]. Nonetheless, the KSVZ and DFSZ models are typically used as benchmarks in axion studies. This is because they represent a large class of models whose couplings to SM particles differ only by factors of order unity. Since the decay constant f_a always remains a free parameter, the condition of the same order of magnitude vaguely defines a band in two-dimensional parameter space. This is often called the QCD axion band [108, 115]. Some effort has been invested into constructing models which lie outside of this band because these can have a crucially different phenomenology and populate experimentally interesting parameter space. Furthermore, it is of theoretical interest to test what kind of properties a QCD axion may have, while still solving the strong CP problem.

An extensive summary of available axion models is given in the recent review [124]. In this reference, methods to enhance or suppress certain couplings are outlined. Such non-canonical axion models can even have a significantly modified relation between m_a and f_a. Reaching larger or smaller masses is difficult to achieve because the additional potential terms should neither spoil the solution of the strong CP problem nor reintroduce fine-tuning. The approaches for building non-canonical axion models are manifold and comprise composite axion models with additional (confined) gauge groups [125–127], large numbers of Higgs doublets [115], clockwork scenarios [128–130] or embeddings into grand unified theories [131–134]. Many attempts have also been made to combine QCD axion models with other BSM scenarios like for instance neutrino masses [24–26], inflation [18–23], gravitational wave production by a first-order phase transition [135–138] or several of these (see e.g. ref. [139] for an attempt to solve five fundamental problems simultaneously).

Explaining all of these models in detail is beyond the scope of this short summary. It is, however, important to note that a large variety of models has been devised, covering a wide range of phenomenologically interesting scenarios and often solving different kinds of problems besides strong CP. Because the suggestions differ so drastically in several important couplings, any detection of an axion will automatically exclude a large number of models.

3.4 Axion effective field theory

The large number of possible UV-complete axion models means that it is inconvenient to individually investigate their low-energy phenomenology in order to search for these signatures. Instead, the method of low-energy effective field theories (EFTs) is perfectly suited to axion physics. This is because the axion is typically the only light new degree of freedom and experiments are probing its interactions to the SM at some scale far below the axion decay constant f_a. In addition, all axion couplings are usually suppressed by f_a and thus the leading order in the EFT expansion will already be sufficient for all practical purposes.[20]

Above the electroweak scale

Before electroweak symmetry breaking, the most general axion EFT including all operators up to mass dimension five can be written as [140–142]

$$\mathcal{L}_{a\text{EFT}} = \mathcal{L}_{\text{SM}} + \mathcal{L}_a \, , \tag{3.19}$$

with

$$\mathcal{L}_a = \frac{1}{2}\partial_\mu a \partial^\mu a - V(a) - \frac{a}{f_a}\sum_F c_{FF}\frac{\alpha_F}{8\pi}F^a_{\mu\nu}\tilde{F}^{a\mu\nu} + \frac{\partial_\mu a}{f_a}\sum_\chi \bar{\chi}C_\chi\gamma^\mu\chi \, . \tag{3.20}$$

The first sum runs over all gauge bosons, $F = B, W, G$, and the indices a are obsolete for the abelian $U(1)_\text{Y}$. In the second sum, χ takes the values of all chiral fermion multiplets in the Standard Model, namely $\chi = Q_L, L_L, u_R, d_R, e_R$. Generation indices have been omitted for simplicity but note that each C_χ is a quadratic hermitian matrix in generation space and allows for flavour-changing fermion currents. If we restrict the EFT to CP conserving interactions[21], which is sensible since the axion is introduced in order to explain the absence of CP violation, the coupling matrices must be symmetric, $c_\chi = c_\chi^T$ (see appendix A). Chirality flipping operators of the form $\frac{a}{f_a}\bar{\chi}\gamma^5\chi$ have not been included because these can be rotated into c_{FF} and C_χ by appropriate chiral rotations of the fermionic fields. Even after this reduction of operators, there is still some degree of redundancy left. Rotations of fermionic fields along baryon number and lepton flavour[22] reduce the number of independent couplings by one in the quark sector and by the number of generations in the lepton sector, resulting in a total of 29 independent real coupling constants for three generations of fermions [142].

After electroweak symmetry breaking

Below the electroweak symmetry breaking scale, fermions as well as the W and Z bosons acquire masses and it is convenient to write the axion interactions in the

---

[20] Note that divergent loop diagrams in the EFT are not necessarily regulated at f_a. Instead many axion EFTs are valid only up to the scale of the lightest BSM particle, e.g. a second Higgs doublet. This is discussed in greater detail in chapter 6.

[21] See ref. [143] for a detailed discussion of CP violating axion interactions in effective field theories, their origin and phenomenology.

[22] Without neutrino masses, lepton flavour is a global symmetry of the SM.

mass eigenbasis. For the gauge bosons this is given by [142]

$$\begin{aligned}\mathcal{L}_a \supset &-\frac{a}{f_a}c_{gg}\frac{\alpha_s}{8\pi}G^a_{\mu\nu}\widetilde{G}^{a\mu\nu}-\frac{a}{f_a}c_{WW}\frac{\alpha_w}{8\pi}W^a_{\mu\nu}\widetilde{W}^{a\mu\nu}-\frac{a}{f_a}c_{ZZ}\frac{\alpha_{\mathrm{EM}}}{8\pi}Z_{\mu\nu}\widetilde{Z}^{\mu\nu}\\ &-\frac{a}{f_a}c_{\gamma Z}\frac{\alpha_{\mathrm{EM}}}{8\pi}F_{\mu\nu}\widetilde{Z}^{\mu\nu}-\frac{a}{f_a}c_{\gamma\gamma}\frac{\alpha_{\mathrm{EM}}}{8\pi}F_{\mu\nu}\widetilde{F}^{\mu\nu}\ .\end{aligned}\tag{3.21}$$

For phenomenological purposes, it is common to define the dimensionful coupling constants

$$\begin{aligned}&g_{agg}=c_{gg}\frac{\alpha_s}{2\pi f_a} \qquad g_{aWW}=c_{WW}\frac{\alpha_w}{2\pi f_a} \qquad g_{aZZ}=c_{ZZ}\frac{\alpha_w}{2\pi f_a}\\ &g_{a\gamma Z}=c_{\gamma Z}\frac{\alpha_w}{2\pi f_a} \qquad g_{a\gamma\gamma}=c_{\gamma\gamma}\frac{\alpha_{\mathrm{EM}}}{2\pi f_a}\ .\end{aligned}\tag{3.22}$$

In the EFT, f_a does not have a direct physical interpretation[23], but only the combination c_{FF}/f_a is observable. One can avoid this redundancy by defining f_a such that $c_{gg}=1$, which is only possible if the axion couples to gluons at all. It appears as if the three independent coupling constants before electroweak symmetry breaking give rise to five couplings after. However, electroweak gauge symmetry implies [142]

$$g_{aWW}=g_{a\gamma\gamma}+\frac{\cos\theta_W}{2\sin\theta_W}g_{a\gamma Z}\ ,\tag{3.23}$$

$$g_{aZZ}=g_{a\gamma\gamma}+\frac{\cos^2\theta_W-\sin^2\theta_W}{2\cos\theta_W\sin\theta_W}g_{a\gamma Z}\ ,\tag{3.24}$$

where the Weinberg angle, $\cos\theta_W=m_W/m_Z$ was introduced. Hence, there are again only three independent couplings to gauge bosons.

Interactions with SM fermions in their mass eigenbasis can be written as

$$\mathcal{L}_a \supset \frac{\partial_\mu a}{f_a}\left(\sum_{f=u,d,e,\nu}\bar{f}_L\,c_{f,L}\,\gamma^\mu\,f_L+\sum_{f=u,d,e}\bar{f}_R\,c_{f,R}\,\gamma^\mu\,f_R\right)\tag{3.25}$$

$$=\frac{\partial_\mu a}{f_a}\left(\sum_{f=u,d,e}\bar{f}\,\left[c_{f,V}\,\gamma^\mu+c_{f,A}\,\gamma^\mu\gamma^5\right]\,f+\bar{\nu}_L\,c_{\nu,L}\,\gamma^\mu\,\nu_L\right)\ ,\tag{3.26}$$

where the first line is in a chiral basis, while the second one splits the interactions into vector and axial-vector types. The coupling matrices are related through $c_{f,V}=(c_{f,R}+c_{f,L})/2$ and $c_{f,A}=(c_{f,R}-c_{f,L})/2$. The seven coupling matrices in this basis are also not independent of each other. Instead, the mass diagonalization involving left-handed doublets results in the relations $c_{\nu,L}=c_{e,L}$ and $c_{d,L}=V^\dagger\,c_{u,L}\,V$ with the CKM matrix V. As a result, the EFTs above and below the electroweak scale have exactly the same number of independent coupling constants.

Below the QCD confinement scale

The next relevant phase transition in the Standard Model takes place at the confinement scale. Quarks hadronise and the new degrees of freedom are mesons and

---

[23] This is apart from quantisation conditions of some axion couplings which are discussed in more detail in chapter 5.

baryons. For the EFT description this has two important consequences. First, the axion-quark terms must be replaced by effective couplings to hadrons. Chiral perturbation theory allows to relate these hadronic couplings to the underlying quark interactions in the UV [54]. For this thesis, the only axion hadron interactions of interest are the ones to nucleons. These can be expressed as

$$\begin{aligned}\mathcal{L}_{aN} &= -ia\,\bar{N}\gamma^5\left(g^0_{aN} + g^3_{aN}\tau^3\right)N &(3.27)\\ &= -ia\,\bar{n}\gamma^5 g_{an} n - ia\,\bar{p}\gamma^5 g_{ap} p\ , &(3.28)\end{aligned}$$

where $N = (p, n)^T$ is the nucleon doublet, g^0_{aN}, g^3_{aN} are the iso-scalar and iso-vector couplings, respectively, and τ^3 is the third Pauli matrix. The second row is expressed in terms of the neutron and proton couplings, which are defined as

$$g_{an} = g^0_{aN} - g^3_{aN} \quad \text{and} \quad g_{ap} = g^0_{aN} + g^3_{aN}\ . \tag{3.29}$$

The reason why these are written as Yukawa-type rather than derivative interactions, like all fermion terms above, is that the calculation of nuclear matrix elements are usually carried out with these types of interactions. Note, however, that much care has to be taken when derivative and Yukawa interactions are related by the equation of motion, as it was done here. Tree-level calculations in each of these bases will agree, but loop corrections, which can be parametrically of the same order as the tree-level result, will not. This issue is discussed in more detail in chapter 5.

After hadronisation, axions can mix with pseudoscalar mesons, most notably the light neutral pions η' and π^0. This will induce a contribution to the axion-photon coupling provided that the axion interacts with gluons at all. At leading order in chiral perturbation theory and assuming that c_{gg} was normalized to unity, this contribution to the photon coupling is given by [54]

$$g^{\text{QCD}}_{a\gamma\gamma} \simeq \frac{\alpha_{\text{EM}}}{2\pi f_a}\left(-\frac{2}{3}\frac{m_u + 4m_d}{m_u + m_d}\right) \simeq \frac{\alpha_{\text{EM}}}{2\pi f_a}\left(-2.03\right)\ . \tag{3.30}$$

The numerical factor in parentheses is called the model-independent contribution to the axion-photon coupling. Its precise calculation beyond the approximation above is rather involved. The latest results are obtained by going next-to-leading order in chiral perturbation theory and using lattice QCD results [54, 69]. By combining the parametrical result in ref. [54] with the updated values and uncertainties from ref. [69] one arrives at

$$g^{\text{QCD}}_{a\gamma\gamma} = \frac{\alpha_{\text{EM}}}{2\pi f_a}\left(-1.93(3)\right)\ . \tag{3.31}$$

All in all, the same EFT description of axion interactions with gauge bosons can be used above and below the confinement scale of QCD. However, sizeable additional contributions from axion meson mixing are expected.

3.5 Axion-like particles

The focus of this thesis are QCD axions, which are defined by their property of solving the strong CP problem via the PQ mechanism. There is however a more general class of hypothetical new particles, called axion-like particles (ALPs). These include pseudoscalars with an approximate shift symmetry which exhibit the same kind of interactions as the QCD axion. They are therefore automatically included in the EFT framework described above. Unlike QCD axions, they do not necessarily solve the strong CP problem because they may not feature the required coupling to gluons or because of other terms in the potential which spoil the PQ solution.

ALPs can emerge from spontaneous breaking of global symmetries or from the compactification of extra dimensions [28, 144, 145]. In particular, many string theory models predict the existence of such particles [144, 144–148]. In fact, they are a common implication of these theories and can even appear in large numbers. This is why the existence of an axiverse was suggested [145, 148, 149], which would be a plenitude of different light and weakly coupled ALPs.

For most phenomenological discussions, it is irrelevant whether a new particle is a QCD axion or an ALP. This is because the observable effects only depend on the low-energy interactions which are often identical. Of course, it would be desirable to confirm the existence of a QCD axion. To this end, a good match of a UV-complete axion model with observations would be a good hint but ultimately a measurement of the axion-gluon coupling would be required.

For the purpose of brevity, the following chapters will generically speak of axions whenever both QCD axions and ALPs are meant to be included. This is also true for the sensitivity studies in chapters 6 and 7, where the EFT parameters are regarded as free. All of the benchmark models in these chapters are indeed QCD axion models.

4 Axion phenomenology and experiments

Parts of this chapter are based on summaries which were published in refs. [1] and [2] before submission of this thesis. Material from these articles is used in the following without further reference.

Already in the first papers proposing the existence of a QCD axion some potential experimental signatures were discussed [12, 13]. The original PQWW axion model, in which f_a is of the order of the electroweak scale, was quickly ruled out in this way [150]. Models with a larger PQ symmetry breaking scale like the KSVZ and DFSZ models were initially called *invisible* because it seemed unrealistic at the time to detect them. However, the progress in experimental techniques and technology has by now certainly proven that this is a misnomer [151]. A growing portion of the most relevant parameter space is coming within reach of current or near future experiments. An extensive overview of the bounds based on these experiments as well as on astrophysical or cosmological observations is given in ref. [152]. It also includes projections of proposed axion searches. In the following, a brief overview of a selection of detection methods is provided.

4.1 Laboratory searches

The most controlled environments for axion detection are pure laboratory searches, where axions are both produced and detected within the experiment. A famous example of such a setup are light-shining-through-wall (LSW) experiments [153, 154]. They rely on the axion-photon coupling $g_{a\gamma\gamma}$. In the presence of a strong magnetic field, axions can convert into a photon via this interaction. This effect can be enhanced by using optical resonators. The resulting axions can freely propagate and leave the resonator, which confines the electromagnetic radiation. Inside a second cavity, these axions can be reconverted into photons via the inverse mechanism and subsequently be detected. The whole procedure corresponds to observing light passing through an opaque wall in the presence of a magnetic field, giving the setup its name. The strongest axion bounds from such an experiment is provided by ALPS [155]. Its update ALPS II is soon expected to start operation [156–158].

LSW experiments quickly lose sensitivity at axion masses $\sim$ meV. This is due to decoherence between the photon and axion waves inside the conversion volume. Significantly heavier axions can be probed in particle colliders. For instance, axions at the heavier end of the currently accessible mass range are best studied at the LHC [141, 159–166] and B-factories [167–173]. For the intermediate regime between MeV and GeV scales, reactors [174], rare decay [64, 167, 168, 171, 172, 175–178], fixed target experiments [179–187] and long-lived particle detectors at the LHC [64, 188, 189] are available (see [64] for a recent review). As discussed in more detail in chapter 6, much lighter axions can also be tested through the $K^+ \to \pi^+ + a$ decay rate at experiments like NA62 [190] and KOTO [191] as well as future experiments such as KLEVER [192].

4.2 Solar axion detection

If axions exist, the Sun must be a powerful source of them close to Earth. The reason is that stars are perfectly suited to the production of light and weakly interacting particles due to their large mass and high temperature. One typical solar axion search is the helioscope concept, devised by Sikivie in 1983 [151]. Its working principle is similar to the one of LSW experiments. An evacuated bore inside a strong magnet is aimed towards the Sun so that the magnetic field is perpendicular to the flux of solar axions. Because the coupling to the electromagnetic field is proportional to $\vec{E} \cdot \vec{B}$, this maximises the probability of an axion converting into a photon inside the magnetic field. The photons carry the same energy as the axions, which puts them in the x-ray spectrum. With appropriate optics, the photons can be focused on an x-ray detector. Any signal that is significantly stronger than the noise level of the detector would indicate the existence of an axion or ALP.

Helioscope searches have the advantage of not requiring axions to contribute to the dark matter abundance in the universe. Instead, the existence of solar axions only relies on appropriate coupling strengths to photons or electrons. In addition, helioscopes are sensitive to a wide axion mass range without the need to search for each possible mass individually as it is often the case in axion DM searches. However, they also lose sensitivity at large masses $\sim 20\,\mathrm{meV}$ because of the same reasons as in LSW. The decoherence can in principle be avoided by using a buffer gas in the conversion volume [193, 194]. This results in an effective mass of the x-ray photons which can restore coherence when it is equal to the axion mass. With a buffer gas, the sensitivity can therefore be improved at high masses at the cost of some finite absorption and the need to scan over different pressure settings.

The helioscope setup also has a few disadvantages. If the axion mass is smaller than $\sim 2\,\mathrm{meV}$ a helioscope is not able to distinguish it from a massless particle [195, 196]. Furthermore, they are not sensitive to QCD axions in the preferred mass range for dark matter. However, with modifications of either the mass-coupling [65–68, 124] or the mass-abundance relation [102–104], an axion that makes up DM could also be detected.

For the phenomenological studies in chapter 7, the probability of an axion converting into a photon inside the helioscope is of crucial importance. The most general result including a buffer gas is given by [197–200]

$$P_{a\to\gamma} = \left(\frac{g_{a\gamma\gamma}B}{2}\right)^2 \frac{1}{q^2 + \Gamma^2/4} \left[1 + e^{-\Gamma L} - 2e^{-\Gamma L/2}\cos\left(qL\right)\right] \,, \tag{4.1}$$

where B denotes the average magnetic field strength perpendicular to the incoming axions, L the length of the bore, Γ the absorption length of x-ray photons in the buffer gas and q the transferred momentum given by

$$q = \frac{m_a^2 - m_\gamma^2}{2\omega} \; . \tag{4.2}$$

m_γ is the effective mass of x-ray photons in the thin buffer gas and ω the energy of the incoming axion. The former is a function of the pressure and can be adjusted to match the mass of the axion. In the vacuum limit, $m_\gamma \to 0$ and $\Gamma \to 0$, the conversion probability takes the form

$$P_{a\to\gamma}^{\text{vacuum}} = \left(\frac{g_{a\gamma\gamma} B L}{2}\right)^2 \times \frac{2(1 - \cos(qL))}{(qL)^2} \; . \tag{4.3}$$

This is further simplified when the axion mass is assumed to be negligible. In this case, $q \to 0$ and the conversion probability becomes energy independent,

$$P_{a\to\gamma}^{\text{massless}} = \left(\frac{g_{a\gamma\gamma} B L}{2}\right)^2 \; . \tag{4.4}$$

All of these results will be used in chapter 7.

So far, the most powerful helioscope is the CAST experiment at CERN [199, 201–208]. It constrained $g_{a\gamma\gamma} < 0.66 \times 10^{-10}\,\text{GeV}^{-1}$ for $m_a < 20\,\text{meV}$ [208]. This is expected to be surpassed in the coming years by the international axion observatory (IAXO). It is a next generation helioscope, which is planned to exceed the sensitivity of all previous broadband axion searches by 1 to 1.5 orders of magnitude in coupling space [200, 209]. The newly probed parameter space will include KSVZ and DFSZ axions in the mass range of 1 meV to 100 meV. Such an improvement becomes possible due to technological advances in superconducting magnets, x-ray optics as well as low-background x-ray detectors. For the first time, the whole setup of a helioscope would be purpose-built including the magnet [210, 211]. In addition, the plan includes mounting the experiment on a rotating platform, allowing Sun tracking for approximately 12 hours per day. There are three design stages to be built one after the other: BabyIAXO, IAXO and IAXO+ each exceeding the sensitivity of the previous one [211]. Details about the different setups including magnet, optics and detectors are discussed in the sensitivity study in section 7.5. Experimental parameters are listed in table 7.5.

Besides helioscopes, which are custom-designed for axion detection, DM direct detection experiments have also performed searches for solar axions and put bounds on their couplings (see ref. [209] for an overview of detection techniques of solar axions based on $g_{a\gamma\gamma}$ or other axion couplings). This has recently received a lot of attention in the context of the observed excess in electron recoil events in XENON1T [212]. However, it was quickly ruled out that the excess is caused by solar axions due to astrophysical constraints [213, 214]. Shortly before submission of this thesis, the first results of the follow-up experiment XENONnT were published. They show no sign of an excess [215].

4.3 Dark matter searches

Besides axion helioscopes, Sikivie also proposed haloscopes as a detection scheme of axion DM in his seminal paper [151]. Haloscopes are currently the most sensitive method to search for axions in the DM window. They consist of a microwave cavity inside a strong magnetic field. When an axion enters the cavity, it can convert into a photon by coupling to the background magnetic field. Because DM is non-relativistic, the frequency of the resulting photon is completely determined by the axion mass. If this frequency is matched by the resonance of the cavity, a strong enhancement of the conversion probability can be achieved. A microwave detector would then be able to detect the induced radiation. Due to the resonance effect, haloscopes can achieve extremely good sensitivities but only in a tiny mass range at a time. Their cavities are constructed in such a way that the resonance frequency can be adjusted. Hence, the accessible mass range of an axion haloscope is defined by the tuning range of its cavity. The strong resonant enhancement comes at the price of having to scan over a multitude of frequency settings in order to cover the axion parameter space. In fact, typical haloscopes only remain in one frequency setting for time periods of the order of minutes while using a high-quality cavity [216]. The scanning results in ragged exclusion lines as depicted in ref. [152].

A large number of axion haloscopes are currently in operation, under construction or have been proposed. A comprehensive list can be found in ref. [152]. Experimental details of the various setups are discussed in ref. [217]. Collectively, these searches will cover or reach into the axion band in the mass range of $0.5\,\mu\text{eV} \lesssim m_a \lesssim \text{meV}$. Note that in order to reach higher axion masses than $\sim 100\,\mu\text{eV}$, one has to go beyond the simplest haloscope design described above. For instance, novel noise reduction techniques [218, 219], alternative resonator designs [220, 221] or broadband searches [222] based on the dish-antenna concept [223] have been suggested. Searching for far below µeV DM axions is yet again a unique experimental challenge, for which the lumped element approach [224, 225] has been proposed.

Finally, axion DM could also be detected via a different coupling than $g_{a\gamma\gamma}$. An example for this kind of search is the CASPEr experiment [226, 227], which employs couplings to gluons or nucleon EDMs. Also conventional direct detection experiments may be sensitive to axions, when they exhibit appropriate couplings to fermions. It was for instance suggested that the XENON1T anomaly [212] could be caused by an axion or axion-like particle [228] even though this would require a suppression of the axion-photon coupling.

4.4 Axions in astrophysics

The core of a star, with typical temperatures of a few keV, a significant density and a large volume, constitutes an ideal source of very light bosons [229–236]. In particular, axions can be produced through their couplings to photons, electrons or nucleons. Once produced, they typically leave the star unimpeded due to their large mean free path.[24] This means that unlike photons, which are only emitted from the surface of an opaque object, axions can be emitted from the entire volume of a star. This mitigates the suppression by small coupling constants. Some of the typical production mechanisms – in particular those that are important in the Sun

---

[24] See ref. [237] for a recent study in which the common assumption of free streaming axions was dropped.

– are discussed in greater detail in chapter 7. If stars like white dwarfs (WD), red-giant branch (RGB) or horizontal branch (HB) stars, supergiants or neutron stars emit axions in sufficient amounts, it can facilitates an additional cooling channel and alter the stellar evolution significantly. The resulting effects can be used to put constraints on the relevant couplings. Summaries of stellar cooling observations used for constraining axions are given in refs. [235, 236].

Additional cooling in WDs is dominated by processes involving electrons and they can therefore only be used to constrain g_{ae}. One observable effect in WDs is the deformation of the luminosity function because WDs will move more quickly towards smaller luminosity if they emit axions in large numbers [238–242]. Furthermore, the rate of the period change of pulsating WDs depends on their cooling efficiency [243–247]. In both of these WD variables, hints for anomalous cooling can be found [239, 240, 244–246]. Moreover, it has been suggested that axions can convert into x-rays after escaping from a magnetic WD and giving an observable signal [248, 249].

Axion emission can also influence the distribution of HB and RGB stars in the Hertzsprung–Russel diagram of globular clusters. For instance, the R-parameter, which is the number density of HB over the one of RGB stars, is a good indicator. An axion coupled to photons would shorten the time stars spend on the HB but hardly influence the cooling of RGB stars. Hence, R would be smaller [250]. Again, constraints [251] and weak hints [252] have been derived from this argument. Another useful observable is the luminosity at the tip of the RGB. Axions coupled to electrons would delay the helium ignition, increasing the luminosity at the RGB tip [253]. Furthermore, in open clusters, the ratio B/R of blue over red supergiants depends on their respective lifetimes. Axion cooling by coupling to photons would significantly shorten only the lifetime of blue supergiants, resulting in a smaller value of B/R [254–257].

Neutron stars can also emit axions in significant numbers from their superfluid core if the coupling to neutrons g_{an} is sufficiently strong. In contrast to all stellar objects mentioned above, the cooling of neutron stars can be observed directly by following the evolution of one star for several years [258]. Again, a slight indication of additional cooling has been discovered [48, 259].

Axions coupled to nucleons can furthermore significantly contribute to the energy loss of supernovae. This would shorten the neutrino burst of such an event. The only supernova whose neutrino burst was observed in several detectors around the world was SN1987A [260–262]. Very soon after this event, the first axion limits were deduced [263]. Several revisions later (e.g. [236, 264, 265]), the very limited available data (~ 24 detected neutrinos [266]) can be used to strongly constrain the coupling strength of an axion to nucleons. With the much improved sensitivity of today's experiments, a supernova similar to SN1987A would be a golden opportunity to derive better constraints or find hints for the existence of axions. Supernovae at much larger distances can provide a limit on the coupling to photons, which is deduced from the expected rate of gamma-rays due to axion decays [267]. Moreover, the combination of all past supernovae can lead to a diffuse flux of highly energetic axions. These could be detectable as gamma-rays when they convert into photons inside the galactic magnetic field [268].

A final astrophysical test and possible hint for the existence of axions is the gamma-ray transparency of the universe. A number of studies have found that gamma-ray sources appear brighter than expected [269–283]. This can be explained by photon-axion oscillations in the intergalactic magnetic fields. High-energy gamma-rays are normally damped by pair production of electrons and positrons. However, in the presence of a magnetic field and if $g_{a\gamma\gamma}$ is large enough, photons can convert to axions, which are not damped as strongly by this effect. These can reconvert to photons after travelling freely for a long distance. The effect is very similar to the one in an LSW experiment but on a cosmic distance scale. Coherent photon-axion oscillations are only possible on these large scales if the axion mass is very small ($m_a \lesssim 10^{-7}\,\mathrm{eV}$) but they still need $g_{a\gamma\gamma} \sim 10^{-11}\,\mathrm{GeV}^{-1}$. This puts these bounds and hints far above the usual QCD axion band. Furthermore, not all authors reach the conclusion of an unexpectedly transparent universe [284–286], thereby only constraining the axion parameter space.

5 Axion periodicity, quantised couplings and loop corrections

This chapter is based on so far unpublished work under the working title *Loop-induced corrections to quantised axion couplings*.

One of the often-neglected properties of a QCD axion is its periodicity, which has been recently pointed out as one of the directions for future studies on axion theory [287]. Only few authors have previously investigated its consequences for low-energy phenomenology [288–290]. The reason for this is that the periodicity only becomes apparent in UV-complete axion models and that it is often not carried along to an effective description in the infrared (IR). Instead, common approximations like the application of the equations of motion at the Lagrangian level or an expansion in small field values, which can in many cases be reasonable simplifications, explicitly break the axion's periodicity. This chapter will emphasize the importance of periodicity conditions, demonstrate how it can help to avoid mistakes in calculating observables and investigate why loop-induced corrections to quantised axion couplings are not in contradiction to the periodicity. Many of the results presented here have already been calculated in other contexts before. They are used here to explain and demonstrate how periodicity conditions are fulfilled in effective field theories.

5.1 Axion as a compact scalar

The QCD axion was introduced above as the pseudo-Goldstone boson of a spontaneously broken global U(1) symmetry. In order to solve the strong CP problem, this symmetry needs to be broken by a chiral anomaly, which generates a coupling to gluons and is the reason for it being a *pseudo*-Goldstone boson.

The concept of periodicity becomes most apparent in the simple case where the axion is just the angular degree of freedom of a complex scalar ϕ with a symmetry breaking potential.[25] When the scalar obtains a VEV, $\langle\phi\rangle = v$, one can write its two real degrees of freedom as

$$\phi(x) = \frac{1}{\sqrt{2}}(v + \rho(x))e^{i\frac{a(x)}{v}} \ . \tag{5.1}$$

By choosing v as the normalization scale in the exponential and $1/\sqrt{2}$ as a prefactor, both the axion field a and the radial degree of freedom ρ automatically obtain

---

[25] More complicated scenarios like the DFSZ model, where the axion is a linear combination of angular degrees of freedom of multiple scalars, are considered in the next section.

a canonically normalized kinetic term without kinetic mixing. Furthermore, it becomes clear that a shift of the axion field by a multiple of $2\pi v$, i.e.

$$a \to a + n\,2\pi v \qquad \text{with} \qquad n \in \mathbb{Z}\,, \tag{5.2}$$

must be a symmetry of the theory, because such a shift leaves the underlying fundamental degree of freedom ϕ invariant [289, 291]. A breaking of this symmetry would render the theory inconsistent and it thus has to be exact. It is in fact a discrete gauge symmetry because such a shift relates two redundant descriptions of the system [291].

The most immediate consequence of this discrete symmetry, which is referred to as the axion's *periodicity*, is that the axion potential as well as all axion interactions must be invariant under these discrete shifts. If the axion were a true Goldstone boson, it would be massless and would only receive derivative interactions automatically fulfilling these conditions. In fact, the only axion interaction which breaks the *continuous* shift symmetry is the anomalous coupling to gluons,

$$\frac{g_s^2}{32\pi^2} N \frac{a}{v} G^a_{\mu\nu} \widetilde{G}^{a\mu\nu}\,. \tag{5.3}$$

N is the anomaly coefficient of the SU(3) chiral anomaly, which was defined in equation (3.3) and is always an integer by construction. Exactly as it was previously discussed in section 2.3 in the context of the axion solution, this interaction breaks the continuous shift symmetry and provides a mass for the axion. The generated potential is however also periodic for topological reasons. It is common to define the axion decay constant f_a such that the axion potential is periodic with period $2\pi f_a$. This leads to

$$f_a = \frac{v}{N} = \frac{v}{N_{\mathrm{DW}}}\,. \tag{5.4}$$

N_{DW} is the domain wall number, which is defined by the relation

$$N_{\mathrm{DW}} = \frac{v}{f_a}\,. \tag{5.5}$$

In the conventions used in this thesis, it turns out to be identical to the anomaly coefficient N.[26] In any case, N_{DW} is always an integer, which ensures that the larger period $2\pi v$ is a multiple of the smaller $2\pi f_a$ and therefore the axion potential always preserves the periodicity.

In summary, one can differentiate between three symmetries of the axion field, which are realised with different levels of accuracy. First, as a pseudo-Goldstone boson the axion has an approximate continuous shift symmetry, which allows to express all couplings to fermions as derivative interactions. This symmetry is not exact but broken anomalously by instanton effects. The resulting potential is periodic with period $2\pi f_a$, which defines a discrete $\mathbb{Z}_{N_{\mathrm{DW}}}$ symmetry. Even though breaking of this symmetry is not discussed in this thesis, it is in principle possible

---

[26] Several authors, like e.g. refs. [108, 124], use a different definition for N, which leads to an integer factor between N and N_{DW}.

(or even generic in general relativity) that this symmetry is broken at a higher scale. Such a breaking can be a solution to the domain wall problem of some axion models because it would allow domain walls to decay [292]. On an even more fundamental level, the axion field is periodic by construction with period $2\pi v = 2\pi N_{\mathrm{DW}} f_a$. This $\mathbb{Z}$ symmetry is the consequence of the axion being a compact scalar by construction and it is thus gauged and exact.

5.2 Periodicity in common models

The connection between the axion's periodicity in an effective theory at low energies and the UV completion is often non-trivial. In this section, the benchmark KSVZ- and DFSZ-type models are investigated. For a more general discussion on how the periodicity of a light axion field emerges in the presence of mixing see ref. [289].

KSVZ-type

In KSVZ-type or hadronic models (see section 3.1) there is just one scalar field charged under the $\mathrm{U}(1)_{\mathrm{PQ}}$ symmetry. The axion is the angular degree of freedom of this scalar. This means that the periodicity simply follows from the polar parametrization as given in equation (5.1). The coupling to $G\widetilde{G}$ is generated by first performing a chiral field redefinition and then integrating out one or more heavy quark fields. In the limit of large fermion masses, its size can be easily derived by calculating the colour anomaly coefficient N [114]. The gauged periodicity is therefore given by $a \to a + 2\pi v$. In the case of $N_{\mathrm{DW}} > 1$, the shift by $2\pi f_a$ is a global symmetry that can in principle be broken without spoiling the theory.

DFSZ-type

The case of DFSZ-type models is substantially more complicated because it includes several pseudoscalar modes and mass mixing terms. Considering the same DFSZ model as in section 3.2 and using the same notation, it is easy to identify the three exact shift symmetries of the three pseudoscalars,

$$a_u \to a_u + 2\pi v_u \quad , \quad a_d \to a_d + 2\pi v_d \quad , \quad a_\Phi \to a_\Phi + 2\pi v_\Phi \ . \tag{5.6}$$

After performing the first orthogonal rotation in the plane spanned by a_u and a_d, one can see that a_w is shifted by $2\pi v_w$ under a shift of either of the two fields.

$$a_u \to a_u + 2\pi v_u \quad \Rightarrow \quad a_w \to a_w + 2\pi v_u \cos\beta = a_w + 2\pi v_w \tag{5.7}$$

$$a_d \to a_d + 2\pi v_d \quad \Rightarrow \quad a_w \to a_w + 2\pi v_d \sin\beta = a_w + 2\pi v_w \tag{5.8}$$

This demonstrates that the new field a_w inherits one gauged shift symmetry from the symmetries of its components. In unitary gauge, the mode a_Z is fixed to zero, which results in the relation $a_u = a_w \cos\beta$. Hence, there are only two shift symmetries left after gauge fixing – one in a_w and one in a_Φ. Because the model also contains mass mixing between these two remaining pseudoscalars, another rotation is required. Note that the potential after spontaneous symmetry breaking, as given in equation (3.12), is symmetric under both $a_w \to a_w + 2\pi v_w$ and $a_\Phi \to a_\Phi + 2\pi v_\Phi$. It is therefore apparent that the model containing a_w and a_Φ (or a_m and a after the rotation to the mass eigenbasis) is symmetric under these two shifts. Typically, however, one is interested in the effective theory containing only the axion a, so the one where the heavy a_m has been integrated out. To leading order, this amounts to fixing the value of a_m to the classical solution at the minimum of the potential,

$$\frac{a_m}{v_m} + \delta = 0 \quad \Rightarrow \quad a_\Phi = a\, \cos\gamma + \text{const}\ . \tag{5.9}$$

The constant only appears for non-vanishing values of δ. After integrating out a_m, there is only one pseudoscalar left – the DFSZ axion. Its period can be directly read off from equation (5.9) by inserting the discrete shift symmetry of a_Φ,

$$a_\Phi \to a_\Phi + 2\pi v_\Phi \quad \Rightarrow \quad a \to a + 2\pi \frac{v_\Phi}{\cos\gamma}\ . \tag{5.10}$$

The fundamental period v of the DFSZ axion, which is related through N_{DW} to the axion decay constant, is therefore given by

$$v = N_{\mathrm{DW}} f_a = \frac{v_\Phi}{\cos\gamma} = \sqrt{(n v_w)^2 + v_\Phi^2}\ . \tag{5.11}$$

The domain wall number can be identified with the anomaly coefficient, which is $3n$ for the DFSZ-type model presented here. This way one arrives at the same expression for f_a as the one given in equation (3.16). This is of course a consequence of the periodicity. Because the axion is periodic with period $2\pi v$, it can only couple to $G\widetilde{G}$ in such a way that the resulting potential preserves this symmetry.

All in all, the effective low-energy theory of a DFSZ axion given in equation (3.18) inherits a discrete exact shift symmetry from the three shift symmetries in the UV (equation (5.6)). This is an explicit demonstration of the fact that a compact field will remain compact after mixing with another massive pseudoscalar or with a scalar which is eaten by a gauge field. These scenarios as well as mixing with non-compact scalars have also been discussed in ref. [289].

5.3 Quantised couplings

Given that the axion is a compact scalar, one can start to investigate the periodicity conditions on effective axion couplings. To this end, it is convenient to consider a toy axion model coupled to just one fermionic field ψ and a gauge field A with field strength tensor F.[27] It is apparent that the axion can couple via derivative interactions just like any shift symmetric Goldstone boson would. In addition, exponential interactions with a discrete shift symmetry are allowed as well as the typical axion-gauge boson interactions. The general Lagrangian for such a model is given

---

[27] For simplicity, all of the calculations in the following sections were performed for a U(1) gauge group. The case of non-abelian groups can however be done completely analogously.

by [30]

$$\begin{aligned}\mathcal{L}_a =& \frac{1}{2}\partial^\mu a \partial_\mu a - \frac{1}{4} F^a_{\mu\nu} F^{a\mu\nu} + (\bar{q}_L i \not{D} q_L + \bar{q}_R i \not{D} q_R) \\ &+ c_1 \frac{\partial_\mu a}{f_a} \bar{q}\gamma^\mu \gamma^5 q - (\bar{q}_L \, m \, q_R e^{i c_2 \frac{a}{f_a}} + \text{h.c.}) \\ &+ c_3 \frac{g^2}{32\pi^2} \frac{a}{f_a} F^a_{\mu\nu} \widetilde{F}^{a\mu\nu} \ .\end{aligned} \tag{5.12}$$

As before, the period of the axion is assumed to be $2\pi v = 2\pi N_{\text{DW}} f_a$. One can start by identifying the conditions imposed by periodicity onto the three coupling constants individually. c_1 is clearly unconstrained since the derivative coupling is shift symmetric. The exponential coupling, on the other hand, is symmetric under a shift by $2\pi f_a / c_2 = 2\pi v / (N_{\text{DW}} c_2)$. Hence, c_2 must be a multiple of N_{DW}^{-1}. This is the simplest example of a quantisation condition for an axion coupling coming from the periodicity of the underlying theory. The case of the c_3 coupling is slightly more complex. Since $F\widetilde{F}$ is a total derivative (see equation (2.10)) one may think that after integration by parts, the interaction is in fact a derivative coupling and therefore unconstrained by periodicity. However, as it was shown in section 2.1, the boundary term can give non-trivial contributions to the action. An example is the usual QCD axion potential generated by instanton effects in the case of an SU(3) gauge theory. Recall that the integral

$$\int \mathrm{d}^4 x \frac{g^2}{32\pi^2} F^a_{\mu\nu} \widetilde{F}^{a\mu\nu} \in \mathbb{Z} \tag{5.13}$$

is a topological winding number and therefore any potential generated by non-perturbative effects must be periodic under shifts by $2\pi f_a / c_3$. Naively, this results in the same quantisation condition for c_3 as for c_2,

$$N_{\text{DW}} c_2 \in \mathbb{Z} \quad \text{and} \quad N_{\text{DW}} c_3 \in \mathbb{Z} \ . \tag{5.14}$$

Note that this argument strictly only applies for couplings to non-abelian gauge fields because no field configuration with finite topological charge exists in abelian theories. However, one can see from the anomaly coefficients (equations (3.1) to (3.4)) that the quantisation of gauge charges in abelian gauge symmetries also leads to quantised axion-couplings. The quantisation of abelian charges in the SM is an observational fact, which could be explained by the existence of magnetic monopoles [293, 294] but they have no been observed so far (see e.g. ref. [295] for a recent negative search result).

The conditions in equation (5.14) have to be fulfilled exactly, when it is assumed that none of the fermions are transforming under the discrete axion shift. While this may seem like an obvious assumption, it is actually very common to perform a non-linear field redefinition of the quark fields which depends on the axion field. This shifts the weights between the different coupling constants. Taking the chiral anomaly into account, the couplings transform under chiral rotations of the fermions as

$$q_L \to e^{i\alpha \frac{a}{f_a}} q_L \quad \text{and} \quad q_R \to e^{-i\alpha \frac{a}{f_a}} q_R \tag{5.15}$$

$$\Rightarrow \quad c_1 \to c_1 - \alpha \quad ; \quad c_2 \to c_2 - 2\alpha \quad ; \quad c_3 \to c_3 + 2\alpha \ . \tag{5.16}$$

Two theories which are related by such a transformation are equivalent. Such a shift does however affect the conditions in equation (5.14). They may not be fulfilled after the chiral rotation because the angle $\alpha \in \mathbb{R}$ can be chosen arbitrarily. However, the sum of c_2 and c_3 is invariant under the field redefinition. Hence, the actual periodicity condition for a general axion theory only implies that

$$N_{\mathrm{DW}}(c_2 + c_3) \in \mathbb{Z} \, , \tag{5.17}$$

which is a weaker condition than the naive one in equation (5.14). Whenever this is fulfilled, a chiral transformation exists after which each term is individually invariant under the axion shift by $2\pi v$. Because there is no other transformation which allows to shift c_2 couplings into c_3, equation (5.17) is the most general quantisation condition that the toy model in equation (5.12) has to obey.

Because of the freedom to perform chiral rotations, there is a redundancy in the operators given in the Lagrangian (5.12). This allows to set one of the three coupling constants to zero. From the perspective of periodicity, one may argue that the most convenient choice is to set c_2 to zero. This has multiple advantages. First, the derivative coupling is periodic independently of the value of the coupling constant leaving c_3 as the only quantised coupling. Second, the exponential interaction consists of an infinite series of operators, which cannot be included in its entirety in a perturbative calculation. That is why it is commonly expanded to linear order as $e^{ic_2 \frac{a}{f_a}} \approx 1 + ic_2 \frac{a}{f_a}$. However, this approximation breaks the periodicity explicitly. It should therefore not come as a surprise that EFTs working with such an operator basis as, e.g. ref. [296], are not periodic in a.

There are two other common approximations in EFTs which break the periodicity explicitly. One is the expansion of the axion potential $V(a)$, which is generated by non-perturbative instanton effects from the axion-gluon coupling. As it was discussed above, $V(a)$ is periodic ($V(a) = V(a + 2\pi f_a)$), but this property is lost when it is expanded around its minimum to $V(a) \approx \frac{1}{2} m_a^2 a^2$. The other periodicity violating approximation is the naive application of the equations of motion at the Lagrangian level after integration by parts. This corresponds to the replacement

$$\frac{\partial_\mu a}{f_a} \bar{q} \gamma^\mu \gamma^5 q \to -2i \frac{a}{f_a} \bar{q} m \gamma^5 q \, , \tag{5.18}$$

which is also equivalent to first doing the chiral transformation but ignoring the anomalous contribution to c_3, followed by an expansion of the exponential. Note that this replacement can also lead to a violation of the quantisation condition.

The axion EFT described in section 3.4 or the one constructed in ref. [142] are examples of how breaking the periodicity by small field expansions or naive application of the equations of motion can be avoided.

5.4 Loop corrections to quantised couplings

So far it was shown that the axion's periodicity is an exact symmetry and that it results in quantisation conditions of certain combinations of axion couplings. However, it is a well-known fact that various axion couplings can be generated or receive corrections at one-loop order when heavier fields or high-energy modes are integrated out [142, 296, 297]. This naively seems to contradict the quantisation conditions because loop corrections are generically continuous functions of the renormalization scale (in the case of UV-divergent terms) or of the masses of particles running in the loop (in the case of finite corrections). This section will address both of these potentially problematic scenarios – the running of axion couplings and finite corrections. It will be shown that the periodicity is conserved when the operator basis is chosen carefully. To simplify the discussion, it is instructive to again consider the toy model given in equation (5.12). Furthermore, one can perform the chiral rotation to a field basis in which $c_2 = 0$.[28] This way, the quantisation condition takes the simple form $c_3 \in \mathbb{Z}$ and one can investigate loop corrections to gauge boson and fermion couplings separately.

5.4.1 Couplings to gauge bosons

In axion-gauge boson interactions, like the one given in equation (5.12), the axion takes the role of a dynamical theta-angle. It is a well-known fact that the operator $g^2 F\widetilde{F}$ in any gauge theory cannot be multiplicatively renormalized [298, 299]. The reason is that the winding number, which is given by the integral over this term, is a topological quantity. This is also why θ is a periodic variable in the first place. One can therefore argue that *because* θ is periodic, the anomalous dimension of the operator $g^2 F\widetilde{F}$ has to vanish and that it need not be subjected to any running under the renormalization group. The same statement holds true when the constant θ is replaced by a dynamical axion field. It is instructive to look at this in two typical choices of operator bases [142, 296],

$$\begin{aligned} \mathcal{L}_{aF} &= c_3 \frac{g^2}{32\pi^2} \frac{a}{f_a} F^a_{\mu\nu} \widetilde{F}^{a\mu\nu} &&= g_{aFF}\, a\, F^a_{\mu\nu} \widetilde{F}^{a\mu\nu} && (5.19) \\ &= c_3\, \boldsymbol{O}_{aF} &&= g_{aFF}\ \tilde{\boldsymbol{O}}_{aF}\ . && (5.20) \end{aligned}$$

Studies of axion effective field theories have shown that the Wilson coefficient g_{aFF} does indeed receive divergent loop corrections [142, 296]. This is in no conflict with the quantisation conditions because it is only $c_3 = g_{aFF} 32\pi^2 f_a / g^2$ which is required to be integer-valued at all scales. Therefore, the quantisation condition ensures that g_{aFF} must exhibit the same renormalization group running as g^2. This way $c_3 \propto g_{aFF}/g^2$ remains constant. This matching of the two beta functions was checked explicitly at one-loop level for SU(3), SU(2) and U(1) gauge theories in ref. [142, 296]. One can formulate this result as a non-renormalization theorem for c_3,

$$\beta_{c_3} = \frac{\mathrm{d}}{\mathrm{d}\log(\mu)} c_3 = 0\ , \tag{5.21}$$

where μ denotes the renormalization scale. This identity is exact to all orders in perturbation theory because it is protected by the axion's periodicity.

---

[28] This is also possible when the full particle content of the SM is taken into account, a fact that is discussed in section 3.4.

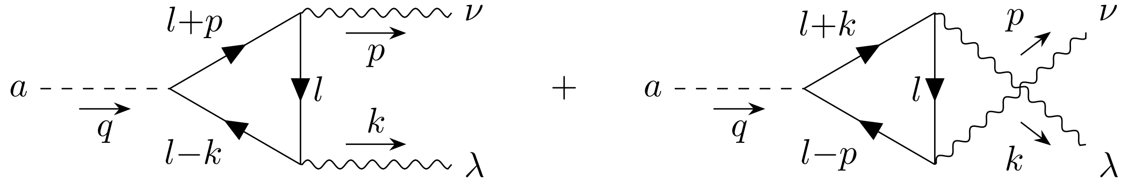

Figure 5.1. Diagrams contributing to the finite corrections of axion couplings to gauge bosons at one-loop level. All Feynman diagrams throughout this work are drawn using the `TikZ-Feynman` [300] package.

It is important to stress that the argument leading to the non-renormalization theorem did not take other axion interactions into account. Instead, it is directly related to the non-renormalization of the theta-angle. When the axion-fermion couplings in equation (5.12) are considered, there are additional one-loop contributions to the gauge boson couplings. The corresponding triangle diagrams are shown in figure 5.1. These are the same diagrams that enter the calculation of the chiral anomaly given in equations (3.1) and (3.2). Several works have calculated these contributions and found them to be finite corrections to the c_3 coupling [142, 296, 297] and they are thus in no conflict with the non-renormalization theorem. Despite this, finite corrections can in principle also violate the quantisation condition when the result – as it is usually done – is interpreted as an effective operator of the type

$$c_3^{\text{eff}} \frac{g^2}{32\pi^2} \frac{a}{f_a} F_{\mu\nu} \widetilde{F}^{\mu\nu} \ . \tag{5.22}$$

This exhibits the core question of this section. How can such a finite correction be in line with the quantisation condition? The most immediate way would be of course if c_3^{eff} was in some way protected by the periodicity, so that it would always turn out to be an integer. This is however not the case. For several fermions with derivative axion coupling $c_{1,f}$, gauge charge Q_f and (colour) multiplicity N_c^f, the on-shell result is given by [141]

$$c_3^{\text{eff}} = c_3 + 4 \sum_f N_c^f Q_f^2 c_{1,f} B(\tau_f) \ . \tag{5.23}$$

The sum runs over all fermions in the theory or just the one fermionic field ψ in the case of the toy model above. The function B depends on the ratio of the fermions and the axion mass, $\tau_f = 4m_f^2/m_a^2$. It is defined as

$$B(\tau) = \begin{cases} 1 - \tau \arcsin^2\left(\frac{1}{\sqrt{\tau}}\right) & \tau \geq 1 \\ 1 - \tau \left(\frac{\pi}{2} + \frac{i}{2} \ln \frac{1+\sqrt{1-\tau}}{1-\sqrt{1-\tau}}\right)^2 & \tau < 1 \end{cases} \ . \tag{5.24}$$

The result given here was calculated in a field basis with $c_{2,f} = 0$, but it was explicitly checked that the same result is recovered after a chiral rotation that instead sets $c_{3,f} = 0$ in the Lagrangian. As a continuous function of the axion and fermion masses, this result will generically not be integer-valued. The questions arise how this result can be reconciled with the axion's periodicity and which step in the calculation resulted in this explicit breaking.

To answer these questions, it is instructive to think about the rationale behind the calculation of c_3^{eff}. In fact, it is a typical effective field theory calculation, where heavy internal fields are integrated out to generate effective interactions between the light degrees of freedom. This will in general result in an infinite tower of higher-order operators. However, the result is clearly a simplification which replaces this infinite series by just one effective interaction term. There are two possible scenarios how this can happen and which could explain the non-integer couplings. Both include explicit breaking of the periodicity, either through a small field expansion or by insertion of the equation of motion.

1. An effective interaction term $aF\widetilde{F}$ could be the lowest dimensional term generated by loop diagrams like the triangles in figure 5.1. The full result would also include terms with an arbitrary power of the axion field a coupled to $F\widetilde{F}$. This can restore the periodicity if the coefficients form a Taylor series of a periodic function. In this scenario, the leading-order term in equation (5.22) would be derived from the full result as
$$\left(c_3^{\mathbb{Z}}\,\frac{a}{v} + c_3^{\mathbb{R}}\sin\left(\frac{a}{v}\right)\right) N_{\text{DW}}\frac{g^2}{32\pi^2}F_{\mu\nu}\widetilde{F}^{\mu\nu} = c_3^{\text{eff}}\frac{g^2}{32\pi^2}\frac{a}{f_a}F_{\mu\nu}\widetilde{F}^{\mu\nu} + \mathcal{O}(a^3)\,. \tag{5.25}$$
On the left-hand side, the first term is a quantised coupling with $c_3^{\mathbb{Z}} \in \mathbb{Z}$. The second coefficient $c_3^{\mathbb{R}}$, on the other hand, can take any real value without conflict with the periodicity, which is only broken by the small field expansion going from left to right. This is similar to the breaking that occurs when the axion potential is approximated to leading order as a mere mass term. The effective coupling is given by the sum of the quantised and the linear order of the periodic function coupled to $F\widetilde{F}$, i.e. $c_3^{\text{eff}} = c_3^{\mathbb{Z}} + c_3^{\mathbb{R}}$.

2. Alternatively, it could be the case that the non-integer result for c_3^{eff} already incorporates an infinite series of higher-order corrections. If each of these higher-order terms is in fact a derivative interaction and thereby conserves the periodicity individually, there would be no contradiction. This scenario can be illustrated as
$$\frac{g^2}{32\pi^2}\sum_{n=0}^{\infty} d_n\left[(\partial_\mu\partial^\mu)^n\,\frac{a}{f_a}\right]F_{\mu\nu}\widetilde{F}^{\mu\nu} \overset{\text{EoM}}{=} \frac{g^2}{32\pi^2}\underbrace{\sum_{n=0}^{\infty} d_n(-m_a^2)^n}_{c_3^{\text{eff}}}\,\frac{a}{f_a}F_{\mu\nu}\widetilde{F}^{\mu\nu}\,, \tag{5.26}$$
where the coefficients d_n have mass dimension $-2n$. In each term the equation of motion, $\partial_\mu\partial^\mu a = -m_a^2 a$, was applied n times. The coefficients will depend on the internal fermion mass and can take arbitrary values except for d_0, which is the only dimensionless and non-derivative coupling. It still has to be integer-valued due to the quantisation condition.

Axion decay in a constant axion background

In order to test whether the first scenario is realised in the toy model introduced in equation (5.12), one needs to calculate the amplitude of n axions to two photons

$(n \cdot a \to \gamma\gamma)$ and check if these form a periodic, analytic function. In classical perturbation theory, it is however not possible to calculate the process with an arbitrary number of final state particles because the possible topologies of contributing Feynman diagrams increases rapidly with n. One can instead take a different approach and calculate the amplitude with just one initial state axion ($n = 1$) but in a constant axion background. This corresponds to the $n \cdot a \to \gamma\gamma$ result in the limit of massless axions and with all axion momenta but one set to zero. This seems like a peculiar limit to take at first sight, but note that if there is an effective interaction of the type $a^n F\widetilde{F}$, it would contribute to the result also in this limit and by a term that would be proportional to the background axion field to the power of $n - 1$. Furthermore, the axion decay rate in an axion background is an observable, which may depend explicitly on the axion field. It therefore has to preserve the periodicity and the following calculation serves as a good check of this principle.

In order to include an axion background, one should work in a field basis in which the axion is coupled in a non-derivative manner, i.e. $c_1 = 0$ and $c_2 \in \mathbb{Z}$,

$$\mathcal{L} \supset -\bar{\psi} m e^{ic_2 \frac{a}{f_a}\gamma^5} \psi \, . \tag{5.27}$$

In a homogeneous axion background, this can be interpreted as an axion-dependent phase of the mass. The resulting fermion propagator is given by

$$\frac{i}{\not{p} - m e^{ic_2 \frac{a}{f_a}\gamma^5}} \, . \tag{5.28}$$

Assuming the standard anti-commutation relations, this can be written as

$$\frac{i}{\not{p} - m e^{ic_2 \frac{a}{f_a}\gamma^5}} = \frac{i(\not{p} + m e^{-ic_2 \frac{a}{f_a}\gamma^5})}{p^2 - m^2} \, . \tag{5.29}$$

Similarly, the axion-fermion vertex also depends on the background axion field and is modified to be

$$\frac{c_2 m}{f_a}\gamma^5 \to \frac{c_2 m}{f_a}\gamma^5 e^{ic_2 \frac{a}{f_a}\gamma^5} \, . \tag{5.30}$$

In order to avoid matrix exponentials in both expressions for the propagator and the vertex, it is helpful to rewrite this as

$$e^{ic_2 \frac{a}{f_a}\gamma^5} = \cos\left(c_2 \frac{a}{f_a}\right) + i\gamma^5 \sin\left(c_2 \frac{a}{f_a}\right) \, . \tag{5.31}$$

It is now straightforward to plug these new Feynman rules into the calculation of the diagram in figure 5.1. The resulting loop integral can be analytically evaluated using the Mathematica program `Package-X` [301]. The result turns out to be independent of the homogenous axion background and it is exactly the same as the one given in equations (5.23) and (5.24).

The calculation proves that the first scenario, where the periodicity is recovered by higher-order non-derivative interactions, is not realised for the loops in figure 5.1. This is not surprising, because the axion background trivially dropped out in the derivative coupling basis and the theory after a chiral rotation is expected to give the same result. Nevertheless, this scenario for periodicity conservation is extremely relevant. For instance, it was shown in appendix A of ref. [288] that it is exactly this kind of effect which reconciles the non-integer coupling generated by axion-pion mixing (see equation (3.31)) with the periodicity condition. The case of axion-pion mixing differs from the mixing with pseudoscalars discussed in the context of the DFSZ model, because in this case the background condition for the heavy field (equation (5.9) in the DFSZ case) is non-linear in the axion field. This results in a series of operators, when the heavy mixing partner like the neutral pion already features a coupling to $F\widetilde{F}$. Interestingly, this also means that the axion decay rate is only independent of a homogenous axion background in the simple toy model above. In typical axion extensions of the SM, where axion-meson mixing is taken into account, the next higher-order operator $a^3F\widetilde{F}$ would also be present and cause a background dependence. Since this effect is suppressed by at least $(a/f_a)^2$ relative to the leading-order decay rate, it may only become relevant when the background field is close to its maximum value, $a \sim f_a$.

Apart from mixing, couplings of periodic functions of a to $F\widetilde{F}$ can also occur as a result of loop corrections. An example for such a situation, which will be further examined in section 5.5, are the effective axion-gluon couplings in refs. [302–304].

Off-shell calculation

One can check whether the effective axion-gauge boson coupling is related to higher-order derivative interactions by avoiding to apply on-shell conditions for the axion. The axion momentum is denoted q as in figure 5.1. If only momentum conservation and on-shell conditions of the photons ($p^2 = k^2 = 0$) are applied, the result for the effective axion-photon coupling reads

$$c_3^{\text{eff}} = c_3 + 4\sum_f N_c^f Q_f^2 c_{1,f} B\left(\frac{4m_f^2}{q^2}\right) , \tag{5.32}$$

with B as defined in equation (5.24). This loop function can now be written as a series in q^2 and it can therefore be identified with an infinite tower of higher-order operators as in equation (5.26). The first term of this expansion vanishes because

$$\lim_{\tau\to\infty} B(\tau) = 0 . \tag{5.33}$$

Therefore, the zeroth order result in axion momenta always obeys the quantisation condition independent of the fermion content of the theory. The fact that the non-integer value of c_3^{eff} is due to effective derivative interactions was recently pointed out in ref. [290].

In summary, two scenarios were outlined which can explain effective non-integer couplings to gauge bosons. Either a truncation of an operator series or the application of the equation of motion is responsible for a result that seems to violate the quantisation and thereby also the axion's periodicity. The latter case can be identified either by making the momentum dependence explicit or by investigating the dependence of the loop result on m_a. In contrast, higher-order interactions which are not proportional to axion momenta will always contribute to observables in a constant axion background. This can also help to distinguish the two cases.

In the toy model of equation (5.12), the effective non-quantised interaction is caused by a derivative coupling evaluated on the equations of motion.[29] The same is true in the Standard Model, where only the diagrams depicted in figure 5.1 contribute at one-loop order [142] to the axion-photon coupling. The first diagrams involving other axion-gauge boson interactions appear at two-loop level. These can in principle result in a periodic function of a coupled to $F\widetilde{F}$ just like the axion's mixings with pseudoscalar mesons. In any case, it is always possible to retrieve an explicitly periodic low-energy Lagrangian by choosing an appropriate operator basis.

5.4.2 Couplings to fermions

Axion-fermion couplings also receive loop contributions in axion EFTs. Because the exponential coupling c_2 is subject to the combined quantisation condition in equation (5.17), it is again non-trivial to understand how the periodicity is conserved when non-discrete radiative corrections are taken into account. To slightly simplify the discussion, it is convenient to start in a field basis where $c_2 = 0$, as it was done in the previous section. By choosing this basis, one only needs to confirm that all effective axion-fermion interactions can either be written as contributions to c_1 or to other higher-order shift invariant interactions so that c_2 is not driven away from zero and the quantisation is left intact. All one-loop processes which couple the axion to two gauge bosons in the toy model of equation (5.12) are depicted in figure 5.2. In typical axion models, there would be a few additional diagrams like the ones labelled 2–4 but with internal scalars instead of gauge bosons or axion-Z mixing contributions [142]. These do not interfere with any of the following arguments.

Since it was assumed that the axion-fermion interactions are of derivative-type at tree level, it is easy to acknowledge that the results of diagrams 2–4 will also be proportional to the axion momentum. Consequently, they only induce contributions to the derivative fermion interactions as computed in ref. [142]. However, when on-shell conditions of the axion are applied, one can again easily transform these to non-quantised pseudoscalar Yukawa-type interactions [296]. Such a result, when used carefully, can still lead to the right phenomenological result. Nevertheless, the correct low-energy EFT *must* include derivative interactions if the periodicity is supposed to be preserved.

The only diagram in figure 5.2 which is not trivially proportional to the axion momentum is diagram 1. It is also the only one which can induce axion-fermion couplings only from tree-level gauge boson interactions. It is therefore the leading-order axion-fermion contribution in axion models without tree-level interactions like the KSVZ model. Its divergent part was first computed by Srednicki [297], but he also interpreted it as an effective Yukawa-type interaction. In order to show that it

---

[29] Note that this is also the explanation for the process dependency of the loop-induced axion-photon coupling described in ref. [305]. The effective axion-photon vertex is not evaluated on the equations of motion when it is used in the Primakoff effect.

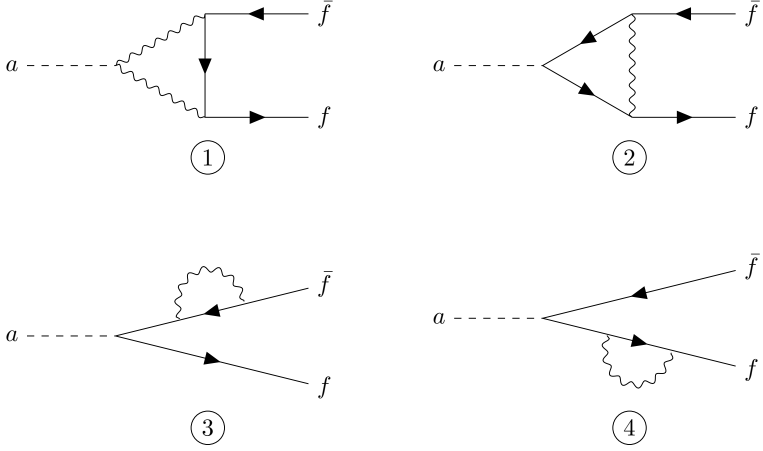

Figure 5.2. Diagrams contributing to the coupling of axions to fermions at one-loop level.

does preserve the periodicity, the calculation can be redone, including both divergent and finite parts and not applying on-shell conditions for the axion. The result can then be written in the compact form,

$$i\mathcal{M}_1 = i\,\bar{v}(p)\,\not{q}\gamma^5\,u(k) \times F_1 \;+\; i\,\bar{v}(p)\,(\not{k}-\not{p})\not{q}\gamma^5\,u(k) \times \frac{m_f}{4m_f^2 - q^2}F_2\;, \tag{5.34}$$

where u and v are the usual positive and negative frequency solutions of the Dirac equation and F_1 and F_2 are dimensionless form factors, which depend on the axion momentum q, the fermion mass m_f and the renormalization scale. UV-divergent terms only appear in F_1. These are responsible for the well-known running of the axion-fermion interactions [142, 296, 297]. F_2 only contains finite corrections, which also preserve the periodicity because the second term can be interpreted as an effective interaction of the type

$$(\partial_\mu a)\,\bar{\psi}(\not{\partial} - \overleftarrow{\not{\partial}})\gamma^\mu\gamma^5\psi\;. \tag{5.35}$$

The arrow indicates a derivative acting on the left-hand side. Clearly, this higher-dimensional operator is symmetric even under continuous shifts of the axion field.

In conclusion, the axion-fermion interactions can be split into those which are subject to quantisation conditions and those which are parametrically free. The exact shift symmetry of the axion field, which is responsible for these conditions, ensures that quantised couplings are not generated radiatively. It is a nice feature of the operator basis chosen in the axion EFT study [142] that this conservation is explicit at all scales. Because some couplings are unrestricted, there can be a renormalization group running of these free parameters. However, the running in theory space is restricted to hypersurfaces which are defined by the quantisation

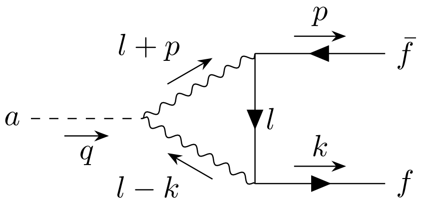

Figure 5.3. Momentum assignments for the calculation of the first diagram in figure 5.2.

conditions (like equation (5.17) in the toy model considered in this section). This way, the periodicity never gets broken explicitly by radiative corrections to axion-fermion couplings.

5.5 Consequences of periodicity

It was shown above that the axion is a compact scalar with a discrete and exact shift symmetry. This symmetry is present at all scales and it imposes periodicity conditions on axion couplings which are not individually shift invariant. This general axion property can also serve as a control mechanism in axion phenomenology. When a new axion effect is investigated – be it in an experiment, cosmology or astrophysics – one can always check whether the result would be altered by a shift in the axion field by $2\pi v$. Often, this will be the case. Then it is important to realise which simplifications were applied to break the periodicity and if the corresponding assumptions are justified in the setup in question. These simplifications will typically be low field expansions or insertion of axion on-shell conditions, i.e. equations of motion. If the restoration of the periodicity by including higher-order terms or replacing Yukawa-type by derivative couplings leads to the same conclusions as the initial discussion, one can be sure that there is no conflict with the axion's periodicity. In the following, two examples will be investigated in this manner. In one case, a non-quantised coupling will turn out to be unproblematic for (almost) all practical purposes, while in the other, the entire model breaks down when higher-order corrections, which are required by periodicity, are included.

Example 1: Axion-pion mixing

An omnipresent effect in axion physics is the model independent contribution to the axion-photon coupling. It appears because axions mix with the lightest pseudoscalar mesons in the confined phase of QCD. This was already mentioned in section 3.4 and it will later become crucial in the context of solar axion detection (section 7.4). By looking at the contribution to c_3 or $g_{a\gamma\gamma}$ given in equation (3.31), it becomes immediately clear that it does not fulfil the quantisation condition, which would be required by periodicity. Exactly as described above, it is helpful to investigate where this apparent contradiction comes from and whether there is a fundamental problem with the calculation. This was also done in appendix A of ref. [288].

To illustrate the effect which leads to the non-quantised coupling, it is sufficient to include mixings with one pseudoscalar meson like the π^0. Its interactions with photons and the combined potential can be written as [288]

$$\mathcal{L}_{\pi a} = \frac{g^2}{32\pi^2}\frac{\pi^0}{f_\pi}F_{\mu\nu}\widetilde{F}^{\mu\nu} + \Lambda^4\cos\left(\frac{\pi^0}{f_\pi}+\frac{a}{f_a}\right) + \Lambda'^4\cos\left(\frac{\pi^0}{f_\pi}\right) \, , \tag{5.36}$$

where Λ and Λ' are energy scales whose value is of minor importance in this context. This potential is clearly periodic in the axion field. Since one is only interested in the effective theory for axions, the heavy π^0 field can be integrated out. The situation therefore seems similar to the mixings in the DFSZ model of sections 3.2 and 5.2. In this case, however, the π^0 is already coupled to $F\widetilde{F}$ and, in addition, the minimisation of the potential leads to a non-linear relation between π^0 and a,

$$\frac{\partial V}{\partial \pi^0} = 0 \tag{5.37}$$

$$\Rightarrow \quad \frac{\pi^0}{f_\pi} = -\arctan\left(\frac{\Lambda^4\sin\left(\frac{a}{f_a}\right)}{\Lambda^4\cos\left(\frac{a}{f_a}\right)+\Lambda'^4}\right) \equiv h\left(\frac{a}{f_a}\right) \, . \tag{5.38}$$

The function h inherits the periodicity of the potential. If it is inserted back into the Lagrangian, the effective axion coupling at all orders in f_a^{-1} reads

$$\mathcal{L}_{a\gamma\gamma\,\mathrm{eff}} = \frac{g^2}{32\pi^2}h\left(\frac{a}{f_a}\right)F_{\mu\nu}\widetilde{F}^{\mu\nu} \tag{5.39}$$

$$= \sum_{n=1}^{\infty} d_n \left(\frac{a}{f_a}\right)^n F_{\mu\nu}\widetilde{F}^{\mu\nu} \, . \tag{5.40}$$

This type of periodicity conservation corresponds to the scenario which was illustrated in equation (5.25). Typically, only the first-order term d_1 is included and therefore the apparent contradiction stems from the small field expansion of h or equivalently from the truncation of the infinite tower of higher order operators.

The question remains, whether the higher-order operators are physically relevant. In most cases, they will clearly be negligible because typical values of the axion decay constant are at least $\gtrsim 10^6$ GeV and thus the next-to-leading order operator will be highly suppressed. However, the effective $3a$-to-2γ vertex also induces a dependence on a possible axion background field. This is similar to the argument made in section 5.4.1. If this background is close to maximal, the axion to photon coupling can significantly deviate from the usual model independent result (equation (3.31)). Such large field values are typical in the early universe, but they occur at temperatures above the QCD phase transition where the quarks do not condensate. One can however envision scenarios in which a strong coupling to matter (for instance neutrons) induces a potential which leads to large field values inside of dense astronomical objects [306]. In such an environment, the higher-order interactions can potentially play a role.

Example 2: Majoron as a QCD axion

In a series of papers [302–304], it was proposed that the QCD axion could be identified with the majoron, which is the pseudoscalar Goldstone boson of a spontaneously broken symmetry generating the neutrino masses. It is helpful to examine this proposal taking the exact axion periodicity into account.

The proposition is intriguingly simple. By construction, the majoron only interacts with neutrinos at tree level. However, it can also couple to gluons via neutrinos, W bosons and quarks. The corresponding three-loop diagram is depicted in figure 5.4. It generates an effective interaction of the type that is required for the PQ solution of the strong CP problem. Thus, the pseudoscalar will acquire a potential which would – as the authors argue – be enough to dynamically set the theta-angle to zero. Interestingly, the result crucially relies on the spontaneous breaking of fermion number conservation and the parity violation in weak interactions [304]. The model predicts an axion symmetry breaking scale (defined as the inverse prefactor of the axion-gluon interaction term), which depends on the weak coupling parameter α_w, the neutrino masses m_ν and another mass scale $\mathcal{M}$, which is of order M_W or the largest heavy neutrino mass [304],

$$f_{a,\text{eff}}^{-1} \propto \frac{\alpha_w^2}{\mathcal{M}^2} \sum_\nu m_\nu \ . \tag{5.41}$$

For reasonable model parameters that are not excluded by observations, $f_{a,\text{eff}}$ is very large ($\sim 10^{18}\,\text{GeV}$ [304]). This is as expected, since f_a is the inverse of a coupling generated only at three-loop order.

Even though this proposed axion model differs from conventional ones in many fundamental ways, the axion (or majoron) is still a compact field with an exact discrete shift symmetry. The reason is exactly the same as given in section 5.1. It is the phase of a fundamental complex scalar and symmetric under discrete shifts by $2\pi v$. As in all the examples which were previously mentioned, the periodicity will be conserved even in the low-energy effective interactions. How exactly this happens in this case is difficult to establish since the three-loop calculations are very involved. Nevertheless, one can easily exclude some possibilities. First, the three-loop result depends on masses and running coupling constants and it is therefore clearly non-quantised. Second, the result does not depend on the axion mass. If it did, this would anyway invert the logic because it is the radiatively generated gluon coupling, which is finally supposed to induce a non-vanishing axion potential. Anyhow, the possibility that the effective $aG\widetilde{G}$ interaction is actually a derivative interaction evaluated on the equations of motion can be excluded. This leaves scenario two from section 5.4.1 as the only viable option: the three-loop result must be a leading term of a periodic function of a/v, where v denotes again the fundamental period of the axion field. Note that this is in no contradiction to any of the calculations in refs. [302–304] but merely relates to the fact that there must be terms of higher order in v^{-1} which lead to the preservation of periodicity.

But what physical consequences do the higher-order terms have? As it turns out they are fatal for this axion model. Similar to what happened with axion-pion mixing, the leading-order term is replaced by some periodic function h,

$$\frac{v}{f_{a,\text{eff}}} \frac{a}{v} G^a_{\mu\nu} \widetilde{G}^{a\mu\nu} \to \frac{v}{f_{a,\text{eff}}} h\left(\frac{a}{v}\right) G^a_{\mu\nu} \widetilde{G}^{a\mu\nu} \ . \tag{5.42}$$

Because the explicit calculation of the $n \cdot a \to g + g$ amplitude at three-loop is beyond the scope of this work, the exact form of h is unknown. Nevertheless, one would

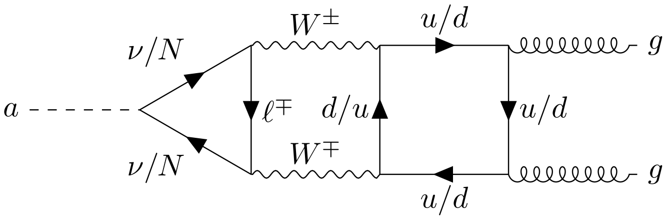

Figure 5.4. Three-loop diagram generating the effective axion-gluon vertex $aG\widetilde{G}$ in the axion model proposed in refs. [302–304]. The coupling is generated via the light SM ν or heavy sterile neutrinos N, charged leptons $\ell^{\pm}$, W bosons and both up- and down-type quarks u/d.

generically expect values of order unity. Since the goal was to device a QCD axion model, this term has to cancel the $\bar{\theta}$ coming from the QCD vacuum structure and complex phases in the quark mass matrix. This leads to the following condition,

$$\left(\bar{\theta} - \frac{v}{f_{a,\mathrm{eff}}} h\left(\frac{a}{v}\right)\right) G^a_{\mu\nu}\widetilde{G}^{a\mu\nu} \overset{!}{=} 0 \tag{5.43}$$

$$\Rightarrow \quad h\left(\frac{a}{v}\right) = \bar{\theta}\,\frac{f_{a,\mathrm{eff}}}{v}\,. \tag{5.44}$$

$f_{a,\mathrm{eff}}$ was found to be larger than v by ~14 orders of magnitude [304]. Because of this and under the assumption $h \lesssim 1$ for all values of the axion field, the equation only has a solution for a if $\bar{\theta}$ is smaller than $\sim 10^{-14}$. Hence, the strong CP problem is not even alleviated in this model. An alternative way to see this, involves the finite field range of the periodic field a. Even if one only took the leading-order term in the effective interaction and interpolated it to very large field values, one would still require a to be of order $f_{a,\mathrm{eff}}$ to cancel theta-angles of order unity. This already contradicts the periodicity when one takes into account that non-compact field ranges should not emerge in the IR [289].[30]

Conclusions

This chapter serves to highlight an important feature of the axion field – its periodicity. It emerges from the UV construction of axion models and is a consequence of it being a Goldstone boson of a spontaneously broken global U(1) symmetry. Neither, mixing with other fields nor loop corrections can interfere with this exact symmetry. The periodicity imposes quantisation conditions on certain axion couplings. The symmetry breaking scale f_a, however, is still arbitrary which means that the couplings measured in experiments like $g_{a\gamma\gamma}$ can still take any value. Nevertheless, it was shown in two examples how the periodicity can help to investigate the phenomenology of certain axion models. In particular, the periodicity becomes an important feature whenever large field values are involved. Furthermore, f_a is directly related to the axion mass and hence the quantisation can be helpful in the

---

[30] A proposal to identify the relaxion as an axion [307] was criticised following similar arguments in ref. [291]. The limited field range was also used in the appendix of ref. [308] to explain why a specific ALP construction does not solve the strong CP problem.

interpretation of experimental data when ratios of quantised couplings and the mass are measured with sufficient accuracy. This is for instance exploited in section 7.4.

6 Leading logarithms in QCD axion effective field theory

This chapter is based on results and arguments which were published under the same title in ref. [2] before submission of this thesis. Material and figures from this article are used in the following without further reference.

EFTs are a regularly applied approach in particle physics phenomenology. The reason is that they allow to test new physics models independent of the details of the underlying model. An EFT can be constructed from a small number of ingredients. First, the field content, i.e. the low-energy fields whose dynamics are described by the EFT, is fixed. In addition, symmetries are imposed, which constrain the set of available operators. The ones of higher mass dimension are suppressed by a typically unknown large energy scale, which may also set a cutoff to the validity of the EFT. Well-known and widely studied examples of such EFTs include Weak Effective Theory [309–311], Standard Model Effective Field Theory [312, 313] and Higgs Effective Field Theory [314–316].

EFTs can also be immensely helpful to perform model-independent phenomenological studies of QCD axions or axion-like particles (see section 3.4). For constructing a minimal axion EFT, one simply imposes a new pseudoscalar particle with a continuous shift symmetry (up to non-perturbative instanton effects). This allows the axion to interact with the SM via dimension-five operator, which either couple it to two gauge bosons or fermions. Such an EFT is very well motivated for the axion because it is thought of as a pseudo-Goldstone boson arising from the spontaneous symmetry breaking of a global and anomalous $\mathrm{U}(1)_{\mathrm{PQ}}$ at a very high scale. This construction both explains the assumed (approximate) shift symmetry, the anomalous gauge interactions and provides a suppression scale for all effective interactions – the axion decay constant f_a. The axion mass is also suppressed by this large scale (see equation (2.28)) and it hence remains very light and can be treated as the only additional low-energy field in the EFT.

Despite this intriguing picture of QCD axion EFTs, it needs to be acknowledged that the adequacy of the EFT language is based on the assumption that the axion is the only manifestation of the Peccei–Quinn mechanism below the scale f_a. In tree-level observables, it is usually already sufficient for all other particles to be much heavier than any of the momenta involved. Loop-level effects, on the other hand, are in principle sensitive to the contributions of heavy off-shell particles running in the loop. These terms can turn out to be the most relevant ones when loop diagrams in the EFT yield divergent contributions. In this case, a cutoff scale has to be included by hand and, as this chapter will demonstrate, it may be more appropriate to identify this scale with the lightest new degrees of freedom in the model (other than the axion) rather than f_a. When constructing an axion EFT from a UV-complete model, it is therefore essential to keep track of the range of validity of the EFT prescription even below f_a.

As an example of such loop-divergent processes, flavour-changing decays involving the axion will be the focus in the following. First, the *leading log* prescription will be introduced, in which the logarithmically divergent term is by far the largest contribution and it is assumed to be cutoff at f_a. This assumption is commonly employed in EFT studies and it can generate large logarithmic enhancements of the amplitudes, thus predicting particularly good sensitivities of future experiments using these observables. However, the following sections will show that typical results in full QCD axion models differ qualitatively and quantitatively from the simplified *leading log* picture because they do not satisfy the implicit assumptions under which the EFT was constructed. This result is in agreement with earlier work in a similar direction [317] where the leading-order renormalization group running was used to demonstrate the absence of large logarithms in flavour-changing effects.

The observation of such a discrepancy between the naive EFT calculation and common axion models leads to the question whether an explicit field theoretic[31] UV-complete QCD axion model exists which truly features a large logarithm in flavour-changing observables. More practically speaking, one would need to generate tree-level couplings of the axion to SM fermions without introducing new degrees of freedom below the PQ symmetry breaking scale. In this thesis, a model of this type is explicitly constructed in two steps. At first, another effective model including the PQ scalar is presented. This model – while not UV-complete – is valid up to $\gtrsim f_a$. It is not entirely new as it coincides with the effective Lagrangian in ref. [317] with specific charge assignments. It also shares some similarity with the flavoured axion models of refs. [23, 318] even though the *leading log* contribution only requires flavour-diagonal axion couplings. For the discussion of leading logarithms, it is anyhow not sufficient to replace one EFT by another. Instead, the higher-energy EFT needs to be UV-completed. This can be done by introducing additional heavy fermionic fields, which serve a similar purpose as the heavy messenger fields in the Froggatt-Nielsen mechanism [319]. The resulting new QCD axion model can then be shown to fulfil all assumptions to reproduce the leading log EFT result for flavour-changing decays. It therefore gives a definitive answer to the question whether UV-complete QCD axion models exist whose EFT description is valid up to the PQ symmetry breaking scale.

The new QCD axion model, which will be introduced below, serves multiple purposes. On the one hand, it proves the theoretical point that the naive EFT result is at least a valid description of some class of QCD axion models. On the other hand, it is a nice benchmark model for axion searches involving flavour-changing effects because it includes all of the required couplings and is rather simplistic at small energy scales. Its phenomenology is interesting and it will be shown that the two upcoming experiments NA62 and IAXO will have good sensitivity to different variations of the model. Despite all of this, the model has some downsides as it requires tuning of the flavour couplings in order to reproduce all SM values without generating large flavour-changing neutral currents (FCNCs). This can be taken

---

[31] In ref. [317] string-motivated UV models were considered for which flavour effects are significantly suppressed.

as further motivation to investigate which axion EFTs have a UV completion and whether some can be disfavoured by introducing mild additional assumptions like the absence of tuning.

6.1 Loop-induced rare meson decays in the EFT

An axion effective field theory for the calculation of meson decay rates can be constructed exactly as described in section 3.4. Since the decays take place at energies far below the electroweak symmetry breaking scale, one is interested in the effective interactions with the mass eigenstates both in the fermionic as well as in the bosonic sector. In principle, the EFT involves all possible effective five-dimensional couplings. For the following discussion however, it is sufficient to consider flavour-universal couplings to SM quarks and leptons.[32] This means that the leading-order flavour-changing neutral current (FCNC) involving axions only appears at one-loop level. Many authors have exploited such loop-induced flavour effects to search for the axion [64, 167, 168, 171, 172, 177, 178, 320, 321].

The assumption of flavour universality at tree level also allows to eliminate any possible vectorial parts of quark and lepton interactions ($c_{f,V}$ in equation (3.26)). This is done by applying flavour-universal vectorial phase rotations, separately of quarks and leptons. Mass terms and charged current interactions with the $W^{\pm}$ are invariant under such transformations and thus no additional Yukawa-like couplings to axions appear. Because of the chirally coupled electroweak fields, there will be an anomalous contribution to axion-gauge boson couplings. This can be absorbed into the corresponding coupling constants, which are anyway present in the EFT. The existence of this contribution to the axion-gauge boson interactions was pointed out in ref. [178]. The flavour-diagonal axion EFT can thus be assumed to only contain the following interaction terms,

$$\begin{aligned}\mathcal{L} \supset\ & \frac{\partial_\mu a}{2f_a}\, c_q \sum_{f=u,d} \bar{f}\gamma^\mu\gamma^5 f + \frac{\partial_\mu a}{2f_a}\, c_l \sum_{f=\ell,\nu} \bar{f}\gamma^\mu\gamma^5 f \\ & - \frac{a}{f_a} c_{WW} \frac{\alpha_w}{8\pi} W^a_{\mu\nu} \widetilde{W}^{a\mu\nu} + \ldots\,, \end{aligned} \tag{6.1}$$

where the neutrino fields are implicitly assumed to be purely left-handed. Both up- and down-type quarks as well as leptons are meant to represent all three generations. The ellipsis stands for other anomalous coupling terms, which are not written out because they are irrelevant for the following calculations.

The minimal axion EFT above can in principle be used to calculate arbitrary flavour-sensitive observables. In order to follow one concrete example and because of its experimental relevance, this chapter focuses on the decay of charged kaons into pions and axions, $K^+ \to \pi^+ + a$. The E787 and E949 experiments at BNL [322] have already constrained the corresponding branching ratio, $\mathrm{Br}(K^+ \to \pi^+ a) < 7.3 \times 10^{-11}$. NA62 expects to improve on this by an order of magnitude by 2025 [323, 324].

---

[32] For axion models with non-trivial flavour properties see refs. [23, 318] or ref. [178] for a summary of the associated phenomenology.

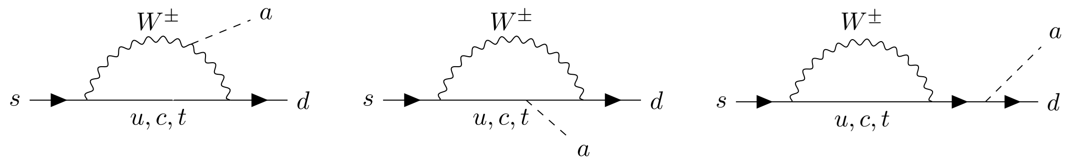

Figure 6.1. One-loop diagrams which induce the flavour-violating $s \to d + a$ transition in the otherwise flavour-universal EFT. The last diagram, a self-energy contribution on the external strange quark leg, arises from the renormalization of quark fields. The analogous one for the down quark line has to be included as well. This equivalence is shown in more detail in appendix D.

Nevertheless, the qualitative results and conclusions are applicable to any FCNC related observable.[33]

On the parton level, the kaon decay is induced by a neutral current of strange and up quarks coupled to an axion, $s \to d + a$. This transition can be parametrized using the effective interaction Hamiltonian

$$\mathcal{H}_{s\to da} = \partial_\mu a\, \bar{d}\gamma^\mu (h^S_{ds} + h^P_{ds}\gamma^5)s + \text{h.c.}\,. \tag{6.2}$$

Working in the usual approximation, where the QCD contributions to the hadronic matrix element can be factored out, one can write

$$\langle \pi(p')|\bar{d}\gamma^\mu(h^S_{ds} + h^P_{ds}\gamma^5)s|K(p)\rangle = h^S_{ds}\, P^\mu f_+(q^2) + h^S_{ds}\, q^\mu f_-(q^2)\,, \tag{6.3}$$

where $q = p - p'$ denotes the momentum transfer and $P = p + p'$. An on-shell axion requires $q^2 = m_a^2$. This means that in the limit of a light axion, the result only depends on the first form factor evaluated at $q^2 \to 0$. Calculations of such form factors require non-perturbative methods. For the analysis in this thesis, a recent lattice evaluation, which arrived at $f_+(0) = 0.9706(27)$, was used. Notably, the matrix element for h^P_{ds} vanishes. This is due to parity, which is (up to possible local effects in quark gluon plasmas [325, 326]) conserved in QCD. Taking all of this into account, the total rate of the decay is given by [317]

$$\Gamma(K^+ \to \pi^+ a) = \frac{|h^S_{ds}|^2}{16\pi m^3_{K^+}}(m^2_{K^+} - m^2_{\pi^+})^2 \lambda^{1/2}(m^2_{K^+},\, m^2_{\pi^+},\, m^2_a)\, f^2_+(m^2_a)\,, \tag{6.4}$$

where $\lambda(x,y,z) = x^2 + y^2 + z^2 - 2(xy + xz + yz)$ denotes the Källén function. Having established this relation between the decay rate and the effective flavour-changing quark interactions, it only remains to calculate h^S_{ds} in the EFT defined in equation (6.1). The types of diagrams which contribute to this transition at one-loop are depicted in figure 6.1. The resulting effective coupling has a logarithmic

---

[33] An example are neutral kaon decays, which were evaluated in refs. [191] and [192].

divergence,

$$\begin{aligned} h^S_{ds} = &- \frac{G_F}{16\sqrt{2}\pi^2}\frac{1}{f_a}c_q \sum_{q=u,c,t} V^*_{qd}V_{qs}m_q^2 \\ &\times \left(\log\left(\frac{\Lambda^2}{m_q^2}\right) - \frac{m_q^4 - 8m_q^2m_W^2 + 7m_W^4 + 6m_W^4\log\left(\frac{m_q^2}{m_W^2}\right)}{2(m_q^2 - m_W^2)^2}\right) \\ &- \frac{3G_F^2m_W^4}{\pi^2}\frac{c_{WW}}{32\pi^2 f_a}\sum_{q=u,c,t} V^*_{qd}V_{qs}f(m_q^2/m_W^2)\,, \end{aligned} \tag{6.5}$$

where G_F is the Fermi coupling constant and V is the CKM matrix. This result agrees with previous calculations [168, 327]. The divergence was regulated by inserting an a priori undetermined UV cutoff Λ. Such a loop calculation can also be understood as the leading-order RG evolution between a high scale Λ and a low scale at which a measurement is performed. The RG approach was chosen for instance in refs. [142, 296, 317, 328]. For convenience, the loop function [168]

$$f(x) = \frac{x(1 + x(\ln(x) - 1))}{(1 - x)^2} \tag{6.6}$$

was used in the W-boson contribution. Furthermore, the result was expanded to leading order in external momenta, i.e. small quark masses. This corresponds to dropping terms which are further suppressed by factors of m_d/m_W, m_s/m_W or m_a/m_W. This way, only the leading-order finite contribution remains. It was also checked explicitly that the result is independent of the chosen field basis and that with a Yukawa-like axion coupling one arrives at the same result as with the derivative coupling when the additional anomalous couplings are taken into account. This is also true in non-unitary gauge where special care has to be taken with the axion Goldstone interactions (see footnote 3 in ref. [175]).

The crucial feature of the effective flavour-changing coupling in equation (6.5) is the logarithmic dependence on an undetermined UV cutoff. It is easy to see that the logarithmic term becomes dominant already for a small hierarchy between λ and the largest quark mass m_t. This is why this term is commonly referred to as the *leading log* and often taken as a good approximation for the full result (see e.g. ref. [327]). Even more intriguingly, the UV cutoff can in principle be larger than m_t by several orders of magnitude. In fact, it is often identified with the only other scale appearing in the EFT Lagrangian – the axion decay constant f_a [171, 175, 327]. Even for a moderate choice of $f_a \sim 10^6\,\text{GeV}$, this results in a logarithmic enhancement of the decay rate by a factor of ~ 300. The UV cutoff is therefore of utmost importance when placing experimental bounds on the EFT parameters in equation (6.1) and, as the next section will demonstrate, the ad hoc choice of $\Lambda \sim f_a$ is typically not justified.

6.2 Comparison to UV-complete models

In order to get a first insight into what determines the correct UV cutoff in the EFT calculation above, it is instructive to look at some explicit examples. Typical UV-complete axion models, which are commonly used to constrain axion parameters, were introduced in chapter 3. In the following, the $K^+ \to \pi^+ + a$ rate will be evaluated in both DFSZ- and KSVZ-type models.

6.2.1 DFSZ-type models

The DFSZ model includes two Higgs doublets in addition to the complex Peccei–Quinn scalar (see section 3.2). This means that the flavour-violating kaon decay is not only generated through the diagrams in figure 6.1. Instead, a large number of different diagrams appear like for instance ones where a W-boson is replaced by a charged Higgs. The results of the full calculation can be taken from ref. [167]. This study investigated the rare B decay $B \to K + a$ but the calculation is identical up to the change of generation indices and hadronic matrix elements. Note, however, that these calculations are performed in a field basis with Yukawa-type axion interactions. As in the EFT case, where it was checked explicitly, this should not affect the result up to terms related to the axion-W coupling.

As outlined in detail in section 3.2, the DFSZ axion inherits couplings to SM quarks from its mixing with the other pseudoscalars in the two Higgs doublet model. Unlike in the EFT in equation (6.5), these are not completely flavour-universal but differ between up- and down-type quarks, $c_u = \cos^2(\beta)/3$ and $c_d = \sin^2(\beta)/3$. Taking into account all diagrams given in ref. [329] the effective coupling of axions to strange and down quarks is given by [167]

$$h^S_{ds} = -\frac{G_F}{16\pi^2}\frac{\cos^2\beta}{3 f_a}\sum_{q=u,c,t} V^*_{qd} V_{qs} m_q^2 \left(X_1^q + X_2^q \cot^2\beta \right). \tag{6.7}$$

As before, the sum runs over all up-type quarks. The structure constants X_i^q are defined as [167]

$$\begin{aligned} X_1^q = {} & 2 + \frac{m^2_{H^\pm}}{m^2_{H^\pm} - m_q^2} - \frac{3m_W^2}{m_q^2 - m_W^2} + \frac{3m_W^4\left(m^2_{H^\pm} + m_W^2 - 2m_q^2\right)}{\left(m^2_{H^\pm} - m_W^2\right)\left(m_q^2 - m_W^2\right)^2}\log\frac{m_q^2}{m_W^2} \\ & + \frac{m^2_{H^\pm}}{m^2_{H^\pm} - m_q^2}\left(\frac{m^2_{H^\pm}}{m^2_{H^\pm} - m_q^2} - \frac{6m_W^2}{m^2_{H^\pm} - m_W^2}\right)\log\frac{m_q^2}{m^2_{H^\pm}}, \end{aligned} \tag{6.8}$$

$$X_2^q = -\frac{2m_q^2}{m^2_{H^\pm} - m_q^2}\left(1 + \frac{m^2_{H^\pm}}{m^2_{H^\pm} - m_q^2}\log\frac{m_q^2}{m^2_{H^\pm}}\right). \tag{6.9}$$

$m_{H^\pm}$ denotes the mass of the charged Higgs boson. The different sign of the leading logarithmic term in the limit of large $m_{H^\pm}$ compared to the result in equation (6.5) is merely caused by a negative constant c_u in the DFSZ model. Plugging this into equation (6.4) allows to compute the decay rate of $K^+ \to \pi^+ + a$. The expected sensitivity of the NA62 experiment to the DFSZ model in terms of the charged Higgs mass and the mixing angle β is plotted in figure 6.2. A remarkable feature is that for each value of $\tan\beta$ there is a value of $m_{H^\pm}$ for which the h^S_{ds} changes sign. This is only possible because X_1^q and X_2^q have a different sign and thus a cancellation appears along a line in figure 6.2 where the sensitivity vanishes completely.

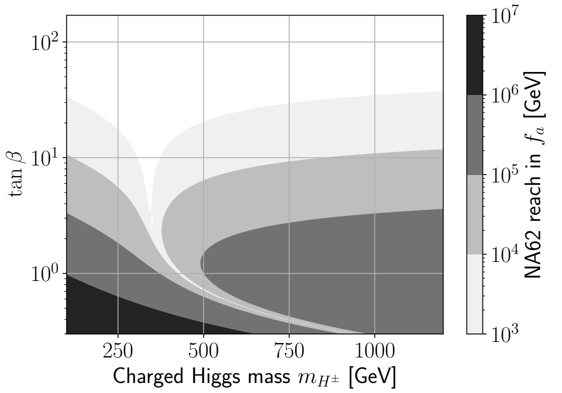

Figure 6.2. Expected sensitivity of NA62 on flavour-universal DFSZ models from $K^+ \to \pi^+ + a$. The reach in f_a is shown in terms of the charged Higgs mass $m_{H^\pm}$ and the mixing angle $\tan\beta = v_u/v_d$. The projection of ref. [323] was used for the numerical evaluation of the sensitivity.

The calculation above explicitly demonstrates that the leading logarithm in effective flavour-changing interactions is not necessarily cutoff at f_a. For the DFSZ model, it is the mass of charged Higgs scalar which regulates the divergence. Since this is close to the electroweak scale, it is orders of magnitude smaller than f_a. But also the cancellations, which can appear for certain parameter combinations, demonstrate the shortcomings of the EFT calculation. Even without a specific combination of parameters, the decay can be suppressed by large values of $\tan\beta$ as becomes clear from figure 6.2. All of this results in a significantly smaller branching ratio of the decay in question.

As a benchmark for the comparison to the EFT calculation and other models, the parameters are chosen as $m_{H^\pm} = 800\,\text{GeV}$ and $1 \leq \tan\beta \leq 5$ (perturbativity requires $0.25 \leq \tan\beta \leq 170$ [124]). On the one hand, this is roughly representative of the experimentally allowed 2HDM parameter space (see e.g. [330, 331]). On the other, it is not a region which is accidentally suppressed by a cancellation between different contributions. The weaker bounds compared to the EFT are therefore predominantly caused by the smaller logarithm. NA62 projections for this benchmark DFSZ model are depicted together with the EFT in figure 6.6.

6.2.2 KSVZ-type models

KSVZ-type models, as introduced in section 3.1, have an apparent advantage over DFSZ-type models – they do not introduce new degrees of freedom below the PQ symmetry breaking scale. This is because the additional fermions are charged under $\text{U}(1)_{\text{PQ}}$ and gain their masses through the spontaneous breaking of this symmetry. Nevertheless, one cannot expect a large, logarithmically enhanced result for the $K^+ \to \pi^+ + a$ branching ratio. This is because of the absence of tree-level couplings to SM fermions in KSVZ-type models. They are only generated at one-loop, an effect which is discussed in section 5.4.2 and which is incapable of generating a large decay rate. Tree-level interactions could otherwise only be introduced by a chiral rotation of SM quarks. But as one would expect, the divergent contribution of a derivative coupling $(\partial_\mu a/(2f_a)\bar{f}\gamma^\mu\gamma^5 f)$ exactly cancels with the one from the perturbatively expanded exponential coupling $(ia/f_a m_f \bar{f}\gamma^5 f)$. This means that only a finite part remains, which depends on the electroweak anomaly coefficients but does not feature a large log.

This argument can be made more explicit and exact for the full six generations of SM quarks, $q = (u, d, s, c, b, t)^T$. Starting without tree-level interactions, the relevant terms for the $s \to d + a$ transition are

$$\mathcal{L} \supset \bar{q} i\gamma^\mu \partial_\mu q - \bar{q} M_q q - \frac{g}{\sqrt{2}} \bar{u}_L \gamma^\mu W^+_\mu V d_L + \text{h.c.}\,. \tag{6.10}$$

A general chiral rotation of each SM quark individually can be written as [141]

$$q \to \exp\left(i\kappa_q \frac{a}{2f_a}\gamma^5 \right) q\,, \tag{6.11}$$

where κ_q is a real 6×6 diagonal matrix. These rotations allow to modify the anomalous couplings to SM gauge bosons. In addition, they generate axion interactions from all of the terms given in equation (6.10). At leading order in the axion field, these are

$$\mathcal{L} \supset -\frac{\partial^\mu a}{2f_a}\bar{q}\kappa_q\gamma_\mu\gamma^5 q - \frac{ia}{f_a}\bar{q}\,\kappa_q M_q \gamma^5\, q - \frac{iag}{\sqrt{8}f_a}\bar{u}_L\gamma^\mu W^+_\mu \left(\kappa_u V - V\kappa_d\right) d_L + \text{h.c.}\,. \tag{6.12}$$

M_q is the diagonalized quark mass matrix and it therefore commutes with κ_q. 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. It becomes clear that the last term only vanishes for completely flavour-universal rotations[34] because in this case $\kappa_u V - V\kappa_d = 0$.

One can collect the contributions of these terms on the $s \to d + a$ transition and it becomes immediately clear that all of the UV-divergent contributions cancel exactly. Thus, the KSVZ model does not contain a large log contribution at one-loop regardless of the chosen field basis.

6.3 A new model inspired by large logarithms

The inability of typical axion models to recover the EFT result with a large UV cutoff begs the question whether any QCD axion model is able to do so. Because an extensive search in axion literature did not come up with such a model, the remainder of this chapter is dedicated to constructing one.

6.3.1 An effective model of the PQ scalar

The guiding principles for constructing a UV-complete QCD axion model with a large logarithmic enhancement in flavour-violating decays can be inferred from the

---

[34] This means that the last term in equation (6.12) appears in the DFSZ model when a basis with derivative axion couplings is chosen. This is because the up- and down-type quarks are not equally coupled to the axion in the Yukawa-basis. It is also the reason why the calculation in the previous section was carried out in the Yukawa coupling basis.

issues of KSVZ and DFSZ models. The previous sections demonstrated that (unlike in the KSVZ) SM fermions must be charged under $\mathrm{U}(1)_{\mathrm{PQ}}$ in order to arrive at tree-level axion interactions with SM quarks. Furthermore, no new degrees of freedom should be introduced below f_a (unlike in the DFSZ) because these may regulate the logarithmic divergence at a lower scale.

It was shown in ref. [118] that any UV-complete QCD axion model must introduce new degrees of freedom beyond the PQ scalar (if one excludes flavoured $\mathrm{U}(1)_{\mathrm{PQ}}$ symmetries). Nevertheless, it is sensible to first search for an EFT only containing SM fields plus the PQ scalar. This would automatically fulfil the assumption of no new degrees of freedom below f_a and could be UV-completed in a second step. A class of such effective models that also satisfy the requirement of tree-level axion-quark couplings are defined by the Lagrangian[35]

$$\mathcal{L} \supset -\frac{\Phi}{\Lambda^u_{ij}}\bar{Q}_{Li}\tilde{H}u_{Rj} - \frac{\Phi}{\Lambda^d_{ij}}\bar{Q}_{Li}Hd_{Rj} + \text{h.c.}\,. \tag{6.13}$$

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

As usual, the PQ scalar Φ contains the axion as its angular degree of freedom. Its PQ charge is normalized to $\chi_\Phi = 1$. The Higgs field is chosen to be uncharged under $\mathrm{U}(1)_{\mathrm{PQ}}$[36] and the quarks charges must therefore satisfy

$$\chi_{Q_{L_i}} - \chi_{u_{R_i}} = \chi_{Q_{L_i}} - \chi_{d_{R_i}} = 1\ . \tag{6.14}$$

The coupling matrices $\Lambda^{u,d}$ have an inverse mass dimension and are thought to be suppressed by a large UV scale. In order to allow all entries of these matrices to be non-zero, the PQ charges are taken to be flavour-independent,

$$\chi_L = \chi_{Q_{Li}} \quad \text{and} \quad \chi_R = \chi_{d_{Ri}} = \chi_{u_{Ri}} \quad \text{for all generations.} \tag{6.15}$$

This is an essential difference to the flavoured axion models which were considered in refs. [23, 318]. The model presented here does not attempt to explain the observed flavour hierarchy of the SM. Nevertheless, flavour effects will become relevant in the UV-complete model.

The QCD anomaly coefficient can be derived from the charge assignments[37],

$$N = \sum_f (\chi_{f_L} - \chi_{f_R}) T(R_f) = 6. \tag{6.16}$$

---

[35] This effective Lagrangian was also considered in ref. [317] based on electroweak symmetry considerations. The authors also provided a string-related axion-like particle realisation to give a large logarithm.

[36] The choice of vanishing Higgs charge under $\mathrm{U}(1)_{\mathrm{PQ}}$ is not a physical property of the model. 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 equation (6.13) 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 and other choices do not alter the physics but merely combine the chiral $\mathrm{U}(1)_{\mathrm{PQ}}$ with the vectorial rotations along baryon number or hypercharge.

[37] Anomaly coefficients differ from the published version in ref. [2] because of a different choice of conventions.

Since it is non-vanishing, the model is indeed able to solve the strong CP problem. The phenomenologically interesting electromagnetic anomaly coefficient

$$E = 2\sum_f (\chi_{f_L} - \chi_{f_R})Q_f^2, \tag{6.17}$$

can either be 16 or 10 depending on the PQ charges of leptons.[38] When the PQ scalar acquires a VEV, $\langle\Phi\rangle = \langle\phi\rangle/\sqrt{2}$, and the heavy radial mode is integrated out, the Lagrangian becomes

$$\mathcal{L} \supset -Y_{ij}^u e^{i\frac{a}{\langle\phi\rangle}} \bar{Q}_{Li}\tilde{H}u_{Rj} - Y_{ij}^d e^{i\frac{a}{\langle\phi\rangle}} \bar{Q}_{Li}Hd_{Rj} + \text{h.c.}, \tag{6.18}$$

where $Y_{ij}^{u,d} = \langle\phi\rangle/(\sqrt{2}\Lambda_{ij}^{u,d})$ are identified as the SM model Yukawa couplings.

After electroweak symmetry breaking, one can rotate into a mass eigenbasis. Because the axion couplings are aligned with the Yukawas, the diagonalization of the quark mass matrix will automatically lead to flavour-universal axion couplings,

$$\mathcal{L} \supset -\sum_q m_q \bar{q}\, e^{i\frac{a}{\langle\phi\rangle}\gamma^5} q. \tag{6.19}$$

This is what was expected from the flavour-independent charge assignments.

All in all, the effective model of the PQ scalar in equation (6.13) combines tree-level couplings to SM fermions with the absence of new particles lighter than f_a. Despite this, the actual value of the UV cutoff is still not defined in this non-renormalizable model. However, one would naively expect a factor of $\log(\langle\phi\rangle^2/m_q^2)$ or $\log(\Lambda_{ij}^2/m_q^2)$ since these are the only scales present in the model. Both are at least as large as the PQ symmetry breaking scale.

6.3.2 UV completion of the model

In order to allow for an explicit calculation of the $s \to d + a$ transition and gain a physical interpretation of the UV cutoff, it is necessary to embed the effective model from the previous section into a renormalizable UV completion. The standard EFT approach in such situations is to interpret the higher-dimensional interactions as a non-local interaction coming from the exchange of a heavy mediator particle. In the case at hand, one could in principle choose to add scalar or fermionic mediators. However, adding a scalar mediator would directly lead back to a two Higgs doublet model, not identical but similar to the DFSZ. This is why the other path is chosen here and three generations of heavy coloured up- and down-type fermions, F_i^u and F_i^d, are added. These are in many ways similar to the messenger fields in the Froggatt-Nielsen mechanism [319]. The approach is illustrated in figure 6.3.

For simplicity and to demonstrate the rationale, the number of generations can be reduced to one. The following calculations are later generalized to include all three generations. The renormalizable interaction terms required for the vertices on the right-hand side of figure 6.3 are given by

$$\begin{aligned} \mathcal{L} \supset \quad & -\alpha^u \bar{Q}_L \tilde{H} F_R^u - \beta^u \bar{F}_L^u \Phi u_R + \text{h.c.} \\ & -\alpha^d \bar{Q}_L H F_R^d - \beta^d \bar{F}_L^d \Phi d_R + \text{h.c.}\ , \end{aligned} \tag{6.20}$$

---

[38] Leptons could also be coupled to Φ^* and carry charges of opposite sign. A number of these kinds of variations of the model are considered in appendix B.

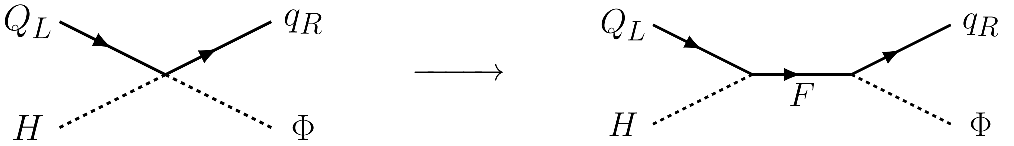

Figure 6.3. The UV completion of the effective model involving only the PQ scalar in equation (6.13) proceeds by introducing new heavy coloured fermions. These are called F quarks. In the limit of small momenta, where they can be integrated out, the effective four particle interaction is recovered.

where $\alpha^{u,d}$ and $\beta^{u,d}$ denote the couplings in the up- and down-type sectors.[39] In principle, one could also consider the possibility to couple either or both of the up- and down-type quarks to Φ^*. This would lead to modified versions of the model, which are described in appendix B. They can be treated in a similar way but this thesis focuses on the version in equation (6.20) for concreteness.

Looking at invariance under the SM gauge groups, it becomes clear that the F^u and F^d fields take the place of right-handed quarks in the SM Yukawa terms. They must therefore live in the respective representations of u_R and d_R. The coupling terms in equation (6.20) also impose conditions on the PQ charges, namely

$$\chi_{Q_L} - \chi_{F_R} = 0 \quad \text{and} \quad \chi_{F_L} - \chi_{q_R} = 1 \,. \tag{6.21}$$

This does not fix the charges precisely, but it is exactly the sum of these equations that appears in the anomaly coefficients N and E in equations (6.16) and (6.17). This is because F and SM quarks couple identically to photons and gluons. As a result, the new fields do not change the anomaly coefficients ($N = 6$ and $E = 16$ or 10) compared to the effective model.

Because the F quarks do not feature chiral interactions in the electroweak sector, gauge invariance allows a bare mass term,

$$\mathcal{L} \supset -\lambda \frac{\langle \Sigma \rangle}{\sqrt{2}} \bar{F}_L F_R + \text{h.c.}. \tag{6.22}$$

The u/d superscript is omitted when referring to both of the two fields. This mass term is parametrized in terms of a VEV of a spurion field Σ which does not carry charge under any SM gauge symmetry or $\mathrm{U}(1)_{\mathrm{PQ}}$. Hence, the charges of F quarks should be independent of their chirality, $\chi_{F_L} - \chi_{F_R} = 0$. This means that all axial charges of fermions, which are the relevant ones for the anomalous couplings to SM gauge bosons, are fixed. The only remaining freedom is a shift of all fermion PQ charges by an arbitrary constant.

The full mass matrix of the combined quark sector appears after PQ and electroweak symmetry breaking. It contains mass mixing terms between the different fields coming from the terms in the Lagrangian (6.20). In unitary gauge, one can

---

[39] The only other dimension-4 term that is invariant under all symmetries considered is $\Phi\Phi^* H H^\dagger$. 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, which is common to all axion models and which is not addressed in this thesis.

define a combined mass matrix M such that all mass terms are included in

$$\mathcal{L} \supset -\begin{pmatrix}\bar{q}_L & \bar{F}_L\end{pmatrix} M \begin{pmatrix} q_R \\ F_R \end{pmatrix} + \text{h.c.}. \tag{6.23}$$

M depends on the VEVs of all scalar fields (H, Φ, Σ), as well as the coupling constants (α, β, λ) and the axion field a, which is the only remaining Goldstone boson from the spontaneous breaking of PQ and electroweak symmetries.

6.3.3 Mass diagonalization

The situation in the new model is in many ways similar to the one known from the SM after electroweak symmetry breaking. The mass matrices M^u and M^d are (mostly) unconstrained complex coupling matrices. In order to recover the quark mass eigenstates, they need to be diagonalized by independent unitary transformations of left- and right-handed fields. In the following, only the main steps are illustrated. A more detailed calculation including all of the relevant intermediate steps is given in appendix C. As before, the u/d labels are dropped whenever the expression is valid for both up- and down-type sectors. Generation indices are also omitted but all three generations of SM and F quarks are meant to be included and the coupling constants α, β and λ are promoted to 3×3 matrices.

The full mass matrix can be written as

$$M = \frac{\langle\Sigma\rangle}{\sqrt{2}} \begin{pmatrix} 0 & \epsilon\epsilon'\alpha \\ \epsilon'\beta e^{ia/\langle\phi\rangle} & \lambda \end{pmatrix}. \tag{6.24}$$

The scale $\langle\Sigma\rangle$ was multiplied out by introducing the two parameters

$$\epsilon = \frac{v}{\langle\phi\rangle} \quad \text{and} \quad \epsilon' = \frac{\langle\phi\rangle}{\langle\Sigma\rangle}, \tag{6.25}$$

with the Higgs VEV $v = 246\,\text{GeV}$. These dimensionless parameters highlight the hierarchies of scales in the model. One can safely assume that ϵ is much smaller than one since the PQ scale must be much larger than the electroweak scale. ϵ', on the other hand, can but does not need to be smaller than one. This depends on the values of the bare mass terms, which are in principle unconstrained. It is only assumed that $\langle\Sigma\rangle$ is larger than $\langle\phi\rangle$ so that the effective theory in equation (6.13) is at least to some precision a low-energy representation of the UV model. Actually, introducing new degrees of freedom very close to $\langle\phi\rangle$ may lead to sizeable corrections like e.g. tree-level flavour-violating axion couplings. In order to avoid such complications, it can be assumed that the additional F quarks are heavy enough so that these effects are sufficiently suppressed. This amounts to the condition

$$\epsilon'' = \frac{\langle\phi\rangle}{\min m_{F_i}} \ll 1, \tag{6.26}$$

where $\min m_{F_i}$ is the mass of the lightest new fermion. The issue of tree-level flavour violation and more concrete bounds on the expansion parameters ϵ and ϵ'' are given in appendix C.

The diagonalization proceeds completely analogously to the SM. One can define the unitary matrix U by its property of diagonalizing the hermitian matrix $MM^\dagger$,

$$U^\dagger M M^\dagger U = \Lambda^2 = \begin{pmatrix} M_q^2 & 0 \\ 0 & M_F^2 \end{pmatrix} . \tag{6.27}$$

M_q and M_F are the diagonal mass matrices of the physical mass eigenstates, i.e. the SM and F quarks. Another unitary matrix is given by $S = M^\dagger U \Lambda^{-1}$. If the left-handed quark fields are now rotated by U and the right-handed ones by S,

$$\begin{pmatrix} q_L \\ F_L \end{pmatrix} \to U \begin{pmatrix} q_L \\ F_L \end{pmatrix} , \quad \begin{pmatrix} q_R \\ F_R \end{pmatrix} \to S \begin{pmatrix} q_R \\ F_R \end{pmatrix} , \tag{6.28}$$

the mass terms in the Lagrangian are diagonalized and the exponential couplings to the axion field a are also eliminated. Conventionally, the same symbols are used for the fields before and after mass diagonalization, despite the fact that they are non-trivial linear combinations of one another.

In order to arrive at an explicit expression of the masses in terms of the coupling matrices in the UV Lagrangian (6.20), one can expand U in ϵ and ϵ''. At leading order, the masses of SM quarks M_q and F quarks M_F are given by

$$M_q^2 = \frac{v^2}{2} \epsilon'^2 U_\delta^\dagger (\alpha \lambda^{-1} \beta \beta^\dagger \lambda^{\dagger -1} \alpha^\dagger) U_\delta , \tag{6.29}$$

$$M_F^2 = \frac{\langle \Sigma \rangle^2}{2} U_\xi^\dagger (\lambda \lambda^\dagger) U_\xi , \tag{6.30}$$

where U_δ and U_ξ are unitary matrices which diagonalize the hermitian matrices in parentheses. In this way, the UV theory defines a total of 12 different quark masses, keeping in mind that both of the expressions above exist for up- and down-type quarks.

Note that any physical choice of UV parameters must be able to reproduce the SM. Equation (6.29) shows that light quark masses are proportional to $\epsilon' v$ times a rather complex combination of coupling constants. One would generically expect this dimensionless factor not to be much larger than one. This hints at a potential problem of the UV-complete model. 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 (which would require $\epsilon' \ll 1$) and when perturbativity of all Yukawa couplings in the UV model is simultaneously required. Similar arguments led the authors of ref. [317] to the conclusion that the cutoff of the effective Lagrangian (6.13) cannot be much higher than f_a. Without fine-tuning of the couplings, this thesis reaches the same conclusion. However, in the explicit UV completion presented here, the masses of the six extra fermions can in principle have a large internal hierarchy. This may still allow enhancements of loop-induced flavour-violating effects by factors larger than $\log(f_a^2/m_q^2)$ and simultaneously generate the large top mass. Some more details on these non-trivial flavour properties of the model are given in appendix C alongside an explicit realisation of coupling matrices for which all of the calculations and conclusions given in this chapter apply.

Axion couplings to gauge bosons

Just like in the KSVZ and DFSZ models described in sections 3.1 and 3.2, the axion couplings to gauge bosons are generated by chiral rotations of fermionic fields. These are included in the transformations by U and S in the previous paragraph because the mass matrix was defined to include the axion field a. To make these field dependent rotations more explicit, it is instructive to split the diagonalization procedure into two parts. First, the axion dependence is absorbed into right-handed quarks and only then the full mass matrix is diagonalized.

Equation (6.20) demonstrates that the dependence on a can be fully absorbed into the right-handed q fields by the redefinitions

$$u_R \to e^{-\frac{ia}{\langle\phi\rangle}} u_R \quad \text{and} \quad d_R \to e^{-\frac{ia}{\langle\phi\rangle}} d_R \,. \tag{6.31}$$

This transformation completely defines the resulting interactions of axions to gauge bosons. In addition, the axion is removed from the mass matrix and instead appears in the kinetic terms of right-handed quarks as a derivative interaction. Therefore, the subsequent steps for fully diagonalizing M can be performed independently of the axion field. The derivative couplings are investigated further in the next paragraph (see equation (6.36)).

Because the path integral measure is not invariant under the chiral redefinition in equation (6.31), anomalous interaction terms are generated just as described in section 2.3. The interactions with the bosonic mass eigenstates of the SM can be found to be [14, 49, 50, 58, 107]

$$\begin{aligned}\mathcal{L} \supset & - N \cdot \frac{\alpha_s}{8\pi \langle\phi\rangle}\, a\, G^a_{\mu\nu}\widetilde{G}^{a\mu\nu} - E \cdot \frac{\alpha_{\text{EM}}}{8\pi\langle\phi\rangle}\, a\, F_{\mu\nu}\widetilde{F}^{\mu\nu} \\ & + E \cdot \frac{\alpha_{\text{EM}}}{4\pi}\frac{s_W}{c_W}\frac{a}{\langle\phi\rangle} F_{\mu\nu}\widetilde{Z}^{\mu\nu} - E \cdot \frac{\alpha_{\text{EM}}}{8\pi}\frac{s_W^2}{c_W^2}\frac{a}{\langle\phi\rangle} Z_{\mu\nu}\widetilde{Z}^{\mu\nu}\,,\end{aligned} \tag{6.32}$$

where s_W and c_W denote the sine and cosine of the Weinberg angle θ_W. Interestingly there are no interactions with the W-boson because only right-handed fields are transformed, which are singlets under SU(2). As in chapter 5, one defines the axion decay constant as

$$f_a = \frac{\langle\phi\rangle}{N_{\text{DW}}}, \tag{6.33}$$

where $N_{\text{DW}} = N = 6$ is the domain wall number or the number of equivalent vacua. Potential consequences of $N_{\text{DW}} \neq 1$ are briefly discussed in section 6.6. With the definition of f_a, the phenomenologically relevant interactions with gauge bosons read

$$\begin{aligned}\mathcal{L} \supset & - \frac{\alpha_s}{8\pi f_a}\, a\, G^a_{\mu\nu}\widetilde{G}^{a\mu\nu} - \frac{E}{N} \cdot \frac{\alpha_{\text{EM}}}{8\pi f_a}\, a\, F_{\mu\nu}\widetilde{F}^{\mu\nu} \\ & + \frac{E}{N} \cdot \frac{\alpha_{\text{EM}}}{4\pi}\frac{s_W}{c_W}\frac{a}{f_a} F_{\mu\nu}\widetilde{Z}^{\mu\nu} - \frac{E}{N} \cdot \frac{\alpha_{\text{EM}}}{8\pi}\frac{s_W^2}{c_W^2}\frac{a}{f_a} Z_{\mu\nu}\widetilde{Z}^{\mu\nu}\,.\end{aligned} \tag{6.34}$$

The EFT parameters, which can be read off directly from this Lagrangian, are listed in table 6.1. If leptons are chosen to also be charged under $\text{U}(1)_{\text{PQ}}$, the axion dependent phase can analogously be absorbed into the right-handed leptons. Such rotations contribute to the E anomaly coefficient giving it the larger value of 16 but the Lagrangian above is still applicable.

Axion-quark couplings

Once the axion field has been absorbed into the right-handed quarks, the mass matrix M is left independent of a. Therefore, the matrices U and S, which are needed for the diagonalization as described in equation (6.28), can only depend on the coupling matrices appearing in the Lagrangians (6.20) and (6.22). Despite the fact that these are unitary transformations, they do not leave the derivative couplings of right-handed quarks invariant because they mix the light PQ charged SM fields with the heavier uncharged F quarks. The detailed derivation of this is outlined in appendix C. At leading order in ϵ and ϵ'', the axion-quark couplings in the mass eigenbasis can be written as

$$\mathcal{L} \supset \frac{\langle\phi\rangle}{2} \begin{pmatrix} \bar{q}_R & \bar{F}_R \end{pmatrix} (\partial\!\!\!/ a)\, \Lambda^{-1} \begin{pmatrix} (\epsilon\epsilon')^2 \mathcal{A}\mathcal{B}\mathcal{A}^\dagger & \epsilon\epsilon'\mathcal{A}\mathcal{B} \\ \epsilon\epsilon'\mathcal{B}\mathcal{A}^\dagger & \mathcal{B} \end{pmatrix} \Lambda^{-1} \begin{pmatrix} q_R \\ F_R \end{pmatrix} \tag{6.35}$$

$$\begin{aligned} =& \frac{1}{\langle\phi\rangle} \bar{q}_R\, (\partial\!\!\!/ a)\, q_R + \frac{\langle\phi\rangle}{2} \bar{F}_R (M_F^{-1} \mathcal{B} M_F^{-1})\, (\partial\!\!\!/ a)\, F_R \\ &+ \frac{v}{2}\, \epsilon'\, \bar{q}_R (M_q^{-1} \mathcal{A}\mathcal{B} M_F^{-1})\, (\partial\!\!\!/ a)\, F_R + \text{h.c.}. \end{aligned} \tag{6.36}$$

$\mathcal{A}$ and $\mathcal{B}$ were defined for notational convenience and are given to leading order in ϵ'' as

$$\mathcal{A} = U_\delta^\dagger \alpha \lambda^{-1} U_\xi \quad \text{and} \quad \mathcal{B} = U_\xi^\dagger \beta \beta^\dagger U_\xi\ . \tag{6.37}$$

It may seem peculiar that the axions only couple to right-handed quarks in this model. This is an explicit basis choice, which will become convenient in the next section. It would also be possible to absorb the axion dependence of the quark mass matrix into both left- and right-handed fields by combining the transformation in equation (6.31) with a vectorial rotation of all quarks. This would immediately generate derivative interactions regardless of chirality.

The benefits of the derivative coupling basis were already discussed in chapter 5 and it is also the most convenient choice for calculating the $s \to d + a$ process. But for recovering the effective couplings in equation (6.19), which the UV completion initially set out to achieve, one has to rewrite the first term in equation (6.36). This can be done by performing a rotation of the right-handed fields (neglecting interactions with the Higgs and CP-conserving anomalous axion terms) and then expanding in small field values

$$\frac{1}{\langle\phi\rangle} \bar{q}_R\, (\partial\!\!\!/ a)\, q_R \to -\frac{a}{\langle\phi\rangle} \bar{q} M_q i\gamma^5 q\, . \tag{6.38}$$

This is exactly the kind of periodicity violating replacement which was discussed in chapter 5 and it should be used with caution and only if all of the underlying assumptions are fulfilled. It nevertheless demonstrates that the model exactly arrives at the anticipated couplings of equation (6.19), which are proportional to the quark masses. In addition, it gives an intuitive reason why these flavour-diagonal derivative couplings exclusively to right-handed quarks do not violate CP. This fact is proven explicitly in appendix A.

The same replacements can be used for the non-universal couplings to F quarks, which leads to

$$\begin{aligned}
&\frac{\langle\phi\rangle}{2}\bar{F}_R(M_F^{-1}\mathcal{B}M_F^{-1})\,(\partial\!\!\!/ a)\,F_R \\
\to -&\frac{\langle\phi\rangle}{2}a\left[\bar{F}\frac{\mathcal{B}M_F^{-1}-M_F^{-1}\mathcal{B}}{2}iF+\bar{F}\frac{\mathcal{B}M_F^{-1}+M_F^{-1}\mathcal{B}}{2}i\gamma^5F\right]\,. \qquad (6.39)
\end{aligned}$$

Hence, their couplings are not proportional to their mass. A result that was expected, given that the UV fields did not feature a chiral charge under $\mathrm{U}(1)_{\mathrm{PQ}}$ and that the mass of the physical fields is mostly generated by the VEV of the uncharged spurion Σ. The fact that the interactions are not flavour-diagonal can lead to CP violation as described in appendix A.

Finally, also the last term of equation (6.36) can be treated in the same manner,

$$\begin{aligned}
&\frac{v}{2}\epsilon'\bar{q}_R(M_q^{-1}\mathcal{A}\mathcal{B}M_F^{-1})\,(\partial\!\!\!/ a)\,F_R \\
\to -&\frac{v}{2}\epsilon' a\left[\bar{q}\frac{\mathcal{A}\mathcal{B}M_F^{-1}-M_q^{-1}\mathcal{A}\mathcal{B}}{2}iF+\bar{q}\frac{\mathcal{A}\mathcal{B}M_F^{-1}+M_q^{-1}\mathcal{A}\mathcal{B}}{2}i\gamma^5F\right]\,. \qquad (6.40)
\end{aligned}$$

Again, there is room for potential CP violating interactions due to flavour non-universality.

In conclusion, the UV-complete model reproduces flavour-universal axion couplings to SM quarks as given in equation (6.38). These simple terms are in contrast to the couplings among F quarks (equation (6.39)) or the ones mixing heavy and light quarks (equation (6.40)), which exhibit a complex flavour structure. Note that this is only true at leading order in ϵ and ϵ''. The latter may not be a good approximation, in which case significant flavour-changing interactions already appear at tree level (see appendix C).

Summary of low-energy couplings

The mass diagonalization procedure laid out in this section leads to a number of low-energy couplings of the physical fermionic mass eigenstates to the new axion. These can be mapped onto the EFT description of section 3.4. All relevant parameters of both models with and without lepton couplings are listed in table 6.1. By construction, the model maps onto a rather simple EFT description at low energies with the axion decay constant being the only free parameter. The details of the UV completion will become important again in the next section, when the explicit cutoff of a divergent process is evaluated in the new model.

6.4 Flavour and CP effects in the new model

6.4.1 Kaon decay rate

The new model introduced in the previous section was inspired by the leading log in the EFT calculation of rare meson decays. It was designed to both feature tree-level interactions with SM fermions as well as no new fields below the PQ breaking

EFT coefficient	UV parameter	EFT model	EFT-ℓ model
f_a	$\langle\phi\rangle/N$	$\langle\phi\rangle/6$	$\langle\phi\rangle/6$
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
$(c_{q,R}-c_{q,L})$	$-(\chi_{q_R}-\chi_{q_L})/N$	$1/6-0$	$1/6-0$
$(c_{\ell,R}-c_{\ell,L})$	$-(\chi_{\ell_R}-\chi_{\ell_L})/N$	0	$1/6-0$

Table 6.1. List of couplings of the new axion model. All parameters are expressed in the EFT basis introduced in section 3.4. Both the general dependence on charges and anomaly coefficients as well as the explicit values for the two models (either with or without tree-level lepton couplings) are given. Note that for all couplings to gauge bosons ($c_{\gamma\gamma}$, $c_{\gamma Z}$ and c_{ZZ}) the interaction terms should be written with α_{EM} as in equation (3.21) and (6.34). The fermion couplings are all understood to be independent of flavour and the $1/6-0$ indicates that only right-handed fields couple at tree level.

scale. This should make the naive estimate of the kaon decay rate in the EFT (equation (6.5)) valid for this model. An explicit calculation including all new degrees of freedom can confirm this claim.

In the previous section, the axion interactions to SM particles were derived. These are important for the phenomenology of the new model but are not sufficient for the unambiguous calculation of UV-divergent processes. Only using the interactions listed in table 6.1 would lead to the same divergences as in the EFT. Only by including the additional vertices induced by the rotation into a mass eigenbasis the divergence can be regulated. All axion couplings to light and heavy quarks were already included in equation (6.36). The calculation will proceed in this derivative basis, where axions only interact with right-handed quarks. In addition, the unitary transformations generate new flavour-dependent interactions of neutral and charged hadronic currents with the weak gauge bosons. The full details and explicit expressions for the unitary transformation matrices U and S are given in appendix C. At leading order in ϵ the charged current interactions can be written as

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

$$\left. + \; W^-_\mu \begin{pmatrix} 0 & 0 & 0 & 0 \\ V^\dagger & 0 & -\epsilon\epsilon' V^\dagger \mathcal{A}_u & 0 \\ 0 & 0 & 0 & 0 \\ -\epsilon\epsilon' \mathcal{A}_d^\dagger V^\dagger & 0 & (\epsilon\epsilon')^2 \mathcal{A}_d^\dagger V^\dagger \mathcal{A}_u & 0 \end{pmatrix} \right] \begin{pmatrix} u_L \\ d_L \\ F^u_L \\ F^d_L \end{pmatrix}, \qquad (6.41)$$

where the coupling matrices $\mathcal{A}_{u/d}$ are the same ones as defined in equation (6.37)

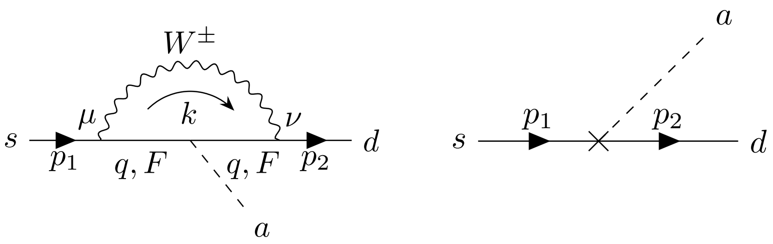

Figure 6.4. One-loop contributions to the $s \to d + a$ process in the new EFT inspired axion model. The left diagram represents four separate contributions with either q or F quarks on both of the internal propagators. The right diagram stands for counterterm contributions from the renormalization of quark fields.

only that it is crucial to differentiate between up- and down-type matrices for the charged current interactions. V denotes the SM 3×3 CKM matrix. At leading order, it can be expressed in terms of the unitary matrices appearing in equation (6.29) as $V = U_\delta^{u\dagger} U_\delta^d$. The first non-unitary term appears at order $(\epsilon\epsilon'')^2$ in the expansion.

Neutral current interactions with the Z boson are also affected by the rotations and come out as

$$\mathcal{L} \supset \begin{pmatrix} \bar{q} \\ \bar{F} \end{pmatrix}^T \gamma_\mu Z^\mu \frac{-g}{\cos\theta_W} \left[\pm\frac{1}{2} \begin{pmatrix} \mathbb{1} & -\epsilon\epsilon'\mathcal{A} \\ -\epsilon\epsilon'\mathcal{A}^\dagger & (\epsilon\epsilon')^2 \mathcal{A}^\dagger\mathcal{A} \end{pmatrix} P_L \right. \\ \left. - Q \sin^2\theta_W \begin{pmatrix} \mathbb{1} & 0 \\ 0 & \mathbb{1} \end{pmatrix} \right] \begin{pmatrix} q \\ F \end{pmatrix}. \tag{6.42}$$

u/d indices are omitted again. The upper and lower signs refer to up- and down-type quarks respectively. P_L denotes the projector onto left-handed fields and Q is the electromagnetic charge of each field. Only the left-handed coupling structure is affected because the heavy F quarks transform differently under SU(2) than the SM quarks but they do carry the same electromagnetic charge.

With these additional vertices at hand, it is finally possible to collect all contribution to the $s \to d + a$ flavour-changing process in the new model. The relevant ones are shown in figure 6.4. The first diagram is meant to include four different contributions because on each internal quark line a light SM or a heavy F quark could be inserted. These diagrams are identical to the ones which were calculated in the EFT, up to different internal masses and more complicated flavour structures at both W and axion vertices as in equations (6.35) and (6.41). Conveniently, the W and axion only interact with left- and right-handed particles, respectively. This means that on each internal quark propagator a mass insertion is required and the momentum term in the numerator drops out. The resulting factors of fermion masses are precisely cancelled by inverse Λ matrices in the axion interaction. Using some of the identities from the diagonalization procedure, it is easy to see that all four diagrams (all left in figure 6.4) have the same flavour structure, are of the same order in ϵ as well as ϵ'' and are parametrically the same in all scales and coupling constants. When the four contributions are added up and the common terms are multiplied out, the only difference between diagrams are sign and masses of internal propagators. The amplitude then takes the relatively compact form

$$\begin{aligned}
&i\mathcal{M}\\
&= i\frac{g^2}{2}(\epsilon\epsilon')^2\frac{\langle\phi\rangle}{2}\sum_{i,j,k,l} V^\dagger_{di}(\mathcal{A}_u)_{ij}(\mathcal{B}_u)_{jk}(\mathcal{A}^\dagger_u)_{kl}V_{ls}\\
&\int\frac{d^4k}{(2\pi)^4}\bar{u}_d(p_2)\left[\gamma_\nu\gamma_\rho(p_2-p_1)^\rho\gamma_\mu P_L\right]u_s(p_1)\\
&\times\left(\frac{1}{k^2-m_W^2}\left(g^{\mu\nu}-\frac{k^\mu k^\nu}{m_W^2}\right)\right)\\
&\times\left(\frac{1}{((p_2-k)^2+m_{Q_i}^2)((p_1-k)^2+m_{Q_l}^2)}-\frac{1}{((p_2-k)^2+m_{F_j}^2)((p_1-k)^2+m_{Q_l}^2)}\right.\\
&\quad\left.-\frac{1}{((p_2-k)^2+m_{Q_i}^2)((p_1-k)^2+m_{F_k}^2)}+\frac{1}{((p_2-k)^2+m_{F_j}^2)((p_1-k)^2+m_{F_k}^2)}\right),
\end{aligned}\tag{6.43}$$

where m_Q and m_F refer to masses of SM and F up-type quarks. The remaining loop integral can be evaluated using `Package-X` [301]. It is helpful to make the simplifying assumption that all F quarks have equal mass. The result at leading order in down and strange quark masses can then be expressed as

$$i\mathcal{M} = h_{ds}\times\bar{u}_d(p_2)\left((m_s-m_d)+(m_s+m_d)\gamma^5\right)u_s(p_1),\tag{6.44}$$

with

$$\begin{aligned}
h_{ds} = &-\frac{G_F}{16\sqrt{2}\pi^2}\frac{1}{\langle\phi\rangle}\sum_{q=u,c,t}V^*_{qd}V_{qs}m_q^2\\
&\times\left(\log\left(\frac{m_F^2}{m_q^2}\right)-\frac{2m_q^4-7m_q^2m_W^2+5m_W^4+3m_W^4\log\left(\frac{m_q^2}{m_W^2}\right)}{(m_q^2-m_W^2)^2}\right).
\end{aligned}\tag{6.45}$$

This is exactly an effective interaction of the type defined in equation (6.4) with $h^S_{ds} = h_{ds}$. One can easily see that the result is qualitatively similar compared to the one of the EFT in equation (6.5). In particular, the leading logarithmically enhanced term is exactly the same but with fixed parameters $c_q = 1/6$ and $\Lambda = m_F$. The finite term was expanded to zeroth order in $1/m_F$ because this is by far the largest mass involved. Despite this, the finite term differs from the EFT result because diagrams involving F quarks also contribute at this lowest order.

In addition to the W boson interactions, the counterterm contributions from the renormalization of quark fields also have to be considered. These are depicted on the right of figure 6.4. The corresponding calculation is included in appendix D and it shows that the relevant contributions cancel exactly at leading order in down and strange quark masses. They can therefore be neglected for all practical purposes.

Since the Z interactions in equation (6.42) are not entirely flavour-universal, one could also consider the first diagram in figure 6.4 with the W replaced by a Z boson.

However, at leading order in ϵ'', which is the limit considered in this section[40], there are no FCNCs between SM quarks induced by Z loops. This is because the coupling matrix structure $\mathcal{A}\mathcal{B}\mathcal{A}^\dagger$ appearing in the relevant amplitudes is identical to the quark mass matrix M_q in equation (6.29) and thereby diagonal.

In conclusion, the new model succeeded in reproducing the large logarithm from the naive EFT calculation. It can therefore be stated with certainty that – even though common QCD axion models are not in agreement with the leading log approximation – at least a class of models exists which possess a UV cutoff larger or equal to the PQ scale.

6.4.2 CP violation

A curious property of the new QCD axion model is the fact that it will generically incorporate additional sources of CP violation. It was discussed in section 3.4 and proven in appendix A that CP violation occurs in flavour-violating interactions if the corresponding coupling constant is complex valued. Because $\mathcal{A}$ or $\mathcal{B}$ in equation (6.36) can have imaginary entries, these conditions are fulfilled in the new UV model even at leading order in ϵ. The reason why this may be problematic is that QCD axion models serve the purpose of explaining the *absence* of CP violation in the strong sector or more concretely the smallness of EDMs. Introducing new CP violating interactions in order to explain CP conservation therefore seems contradictory at first. However, if the additional CP violation does not spoil the PQ mechanism and does not lead to EDMs which are in conflict with observations, one would still have a good QCD axion. CP violation can even be beneficial because it offers new discovery opportunities [332, 333].

To check that the CP violation in the new model is indeed not problematic, it is first important to note its connection to the CP-odd phase of the CKM matrix. Both break CP explicitly and are not due to spontaneous symmetry breaking. In fact, 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, U_δ and U_ξ appearing in equations (6.29) and (6.30) can be chosen to be orthogonal and real if the matrices which 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. In this scenario, no CP violating interactions appear in equation (6.36). The CP violation in the model is therefore already explicitly encoded in the complex UV interactions of equation (6.20). Furthermore, the CKM matrix V is also constructed from these matrices and it is experimentally confirmed to be complex valued. Hence, there is no reason to expect $\mathcal{A}$ and $\mathcal{B}$ to be real without a fine-tuned cancellation.

It remains to be shown that the CP violating axion interactions are not in conflict with EDM observations just like the CKM phase in the SM. In the following, only the neutron EDM d_n is considered because it puts the most stringent bounds on CP violation. d_n can be related to the EDM of free quarks [334] and it vanishes in the limit of no quark EDMs. It was shown in appendix E of ref. [2] that the new model does not generate these quark EDMs at one-loop, which means the

---

[40] See appendix C for consequences of large ϵ''.

leading contributions can only appear at two-loop level. At this point, it is sufficient to conservatively estimate the size of these effects instead of delving into involved two-loop calculations. EDM diagrams must involve the axion since its couplings are the source of CP violation. Therefore at least two new physics vertices with a suppression by f_a must appear. This allows to estimate the order of magnitude of the induced dipole moment as

$$d_n^{\mathrm{UV}} \lesssim \frac{e}{(16\pi^2)^2} \frac{m_n}{f_a^2} \approx 5 \times 10^{-19}\, e\, \mathrm{fm}\,. \tag{6.46}$$

m_n is the neutron mass, which is required in order to fix the dimensions, but it is also the characteristic scale of the system in question. The result is a couple of orders of magnitude smaller than the experimental bound in equation (2.13). The axion decay constant was taken as $f_a = 4 \times 10^6\,\mathrm{GeV}$. This is a conservatively small value considering the range of interest as shown in figures 6.5 and 6.6. Furthermore, the estimate is even more conservative because it does not include any suppression due to new physics contributions by $W^\pm$ or Z gauge boson interactions or the heavy F quark mass scale $\langle \Sigma \rangle$ even though at least one of the three has to appear in order to generate an EDM. Suppressions by light quark masses or their small Yukawa couplings are also not included.

Additional contributions to the neutron EDM may come from interactions between the constituent quarks. This is for instance discussed in the review [62]. These effects are already far below the experimental bounds in the SM and they are only altered in the UV model via terms suppressed by the new physics scale f_a. This leaves them completely negligible.

In conclusion, the new model introduces additional sources of CP violation in comparison to the SM. These are however not in conflict with current observations and it is even unlikely that its CP effects would be detectable in near future experiments. The solution of the strong CP problem is not spoiled and the CP odd couplings still allow the new model to be a proper QCD axion.

6.5 Discovery opportunities

The new QCD axion model introduced in the previous sections was designed to serve a specific purpose. It demonstrates that a UV-complete axion model exists which maps onto the minimal effective model in section 6.1 including the large log enhancement in flavour observables appearing at one-loop level. Its phenomenology could therefore be anticipated from the onset: Tree-level fermion interactions are particularly relevant for axion production in astrophysical objects and the leading log enhancement increases the sensitivity of experiments looking for rare meson decays. In the following, both of these detection opportunities are discussed in more detail.

6.5.1 Astrophysical limits and future searches

For the same reasons as outlined in section 4.4, the new axion can be produced in large numbers inside of astrophysical objects like the Sun, supernovae, horizontal

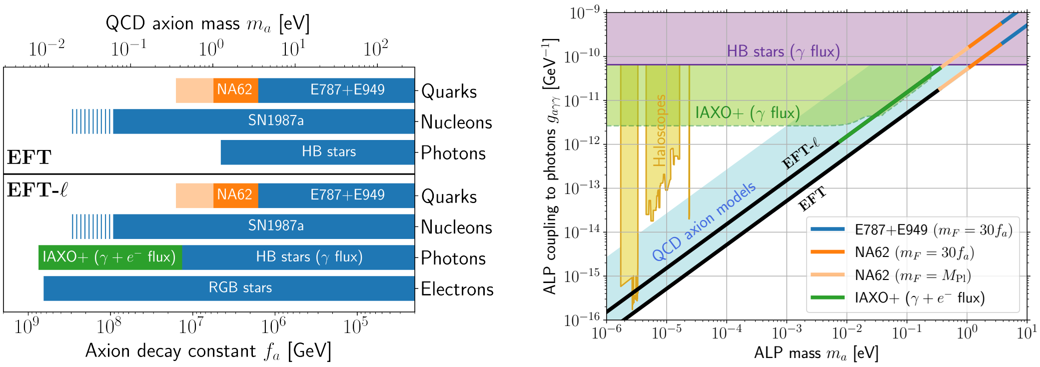

Figure 6.5. Bounds on the two models introduced in section 6.3. They are labelled EFT or EFT-ℓ depending on whether tree-level axion-lepton couplings are included or not. Predicted sensitivities of NA62 and IAXO are shown in orange and green, respectively. The size of the flavour-changing quark couplings depends on the UV cutoff. The lighter shade of orange therefore represents the range of possible values of the cutoff scale for the leading log, $\Lambda \in (30 f_a,\, M_{\text{Pl}})$. Both plots demonstrate that IAXO becomes particularly sensitive to the EFT-ℓ model due to a larger photon coupling and solar axion production through axion-electron interactions. In the left panel, existing constraints on the only free parameter f_a are shown in blue. The hatched region in the supernova bounds on the nucleon coupling indicates the large uncertainties related to this specific observation. On the right, bounds and sensitivities in the mass-photon coupling-plane are shown. The QCD axion band corresponds to the axion window $E/N \in (44/3,\, 5/3)$ as defined in ref. [115].

branch stars or red giants. For the calculation of all of these effects, the tree-level interactions of the axion with SM particles are sufficient. The astrophysical bounds on the new model can therefore be directly inferred by comparing the low-energy effective couplings listed in table 6.1 with the latest axion limits. For instance, the latest helioscope search performed by CAST [335] together with the absence of exotic cooling in horizontal branch (HB) stars in globular clusters [251], constrain the photon coupling to be $g_{a\gamma\gamma} \lesssim 7 \times 10^{-11}\,\text{GeV}^{-1}$. Using the respective values for E/N in the two models, this translates into the bounds on f_a shown in figure 6.5. In the EFT-like model without lepton coupling, there is an accidental cancellation between E/N and the model-independent contribution to the photon coupling (see equation (3.31)). This means that the bound on f_a related to the photon coupling is much stronger for the EFT-ℓ model, which includes tree-level lepton interactions. The upcoming helioscope IAXO [211] is expected to improve the limit on $g_{a\gamma\gamma}$ by more than an order of magnitude. As shown in figure 6.5, this will make it sensitive to the EFT-ℓ model, but it is not sufficient for reaching into the parameter space of the other EFT-like model. This is due to a combination of said cancellation, the signal suppression by larger axion masses and the absence of electron interactions. Such couplings would open several new production channels for axions inside the Sun. These are discussed in greater detail in chapter 7.

Another detection technique relying on the photon interaction are haloscopes (see section 4.3). They have an excellent reach in the mass range of $1 - 100\,\mu\text{eV}$ if the axion makes up the entire dark matter of the universe. Figure 6.4 shows the currently excluded regions as given in refs. [64, 118] in combination with the latest ADMX data [336, 337]. Future searches will keep probing larger parts of the parameter space of both EFT and EFT-ℓ models.[41] Again, the larger photon interaction of the EFT-ℓ model means that it is more easily detectable.

Astrophysical constraints also arise from the effective coupling of axions to nucleons. These are a low-energy manifestation of interactions with gluons and quarks above the confinement scale. It is usually parametrised in the same fashion as other fermion couplings as

$$\mathcal{L} \supset \frac{\partial_\mu a}{2f_a} \sum_{N=p,n} c_N \, \bar{N}\gamma^\mu\gamma^5 N \,, \tag{6.47}$$

where p and n stand for protons and neutrons, respectively. The individual quarks couplings can be matched to specific values of the c_N [54, 338]. In the new axion model, this results in $c_p = -0.39875$ and $c_n = 0.05125$. The strongest bounds on these couplings are derived from the formation of a proto-neutron star in the process of a core-collapse supernova [233]. All limits are based on the observation of the supernova 1987A [260–262]. The most recent ones are given in refs. [265, 339, 340]. In figure 6.5 the spread of these various evaluations is indicated by the hatched regions.

Finally, astrophysical observations also put bounds on the axion coupling to electrons. At tree level, this interaction only exists in the EFT-ℓ model which means that all bounds on the other model are suppressed by a loop factor [124, 297] of $\sim \alpha_{\text{EM}}^2/(\pi^2)\log \sim 10^{-4}$. This makes these bounds irrelevant for any model with $c_e = 0$ in the UV. The strongest limits on the coupling are derived from the brightness of the tip of the red-giant branch (RGB) in globular clusters. These exclude $f_a/c_e \geq 1.9 \times 10^9$ GeV [242]. Other limits of this coupling are based on the R parameter in globular clusters [341, 342]. As figure 6.5 illustrates, the electron coupling puts the strongest bounds on f_a when tree-level axion couplings are present.

6.5.2 Detection via rare kaon decays

By design, one of the most important phenomenological features of the new axion model is the presence of a large logarithmic enhancement in the loop-induced decay of mesons. This results in good sensitivities of fixed targets experiments looking for this type of decays. The flavour-changing process with the best experimental perspectives is the $s \to d + a$ transition.

The predicted decay rate has been computed in equation (6.44). Unlike all the tree-level couplings, which only depend on the single free parameter f_a, the result for h_{ds} also changes with m_F, which is a property of the UV-complete theory. In principle, the only requirement is that m_F must be larger than the axion decay constant. But because the SM Yukawa couplings are also proportional to $\langle\phi\rangle / m_F$ (i.e. they are suppressed by ϵ'), one can derive a rough upper bound on m_F by demanding perturbativity of Yukawa couplings. At least one of the F quarks must therefore have a mass $m_{F_j} \lesssim 10^3 f_a$. Nevertheless, other F quarks may be significantly heavier

---

[41] An extensive overview of existing bounds and sensitivity projections of planned experiments is given in the ref. [152].

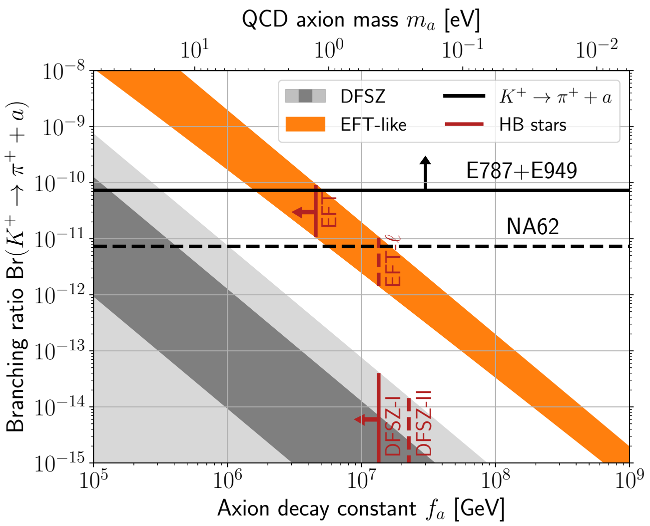

Figure 6.6. Branching ratio of $K^+ \to \pi^+ + a$ predicted in the EFT-like (orange band) and the DFSZ model (grey band). The cutoff scale of the EFT-like model is chosen to be between $30f_a$ and M_{Pl} as discussed in the text. For the DFSZ bands, the charged Higgs mass was fixed to the benchmark value of $m_{H^\pm} = 800\,\text{GeV}$. Typical values of the mixing angle ($1 < \tan\beta < 5$) are plotted in dark grey while the full range allowed by perturbativity ($0.25 < \tan\beta < 170$) is shown in a lighter tone. The horizontal black lines mark the existing constraints by the E787 and E949 experiments [322] (solid line) and the expected reach of NA62 [323] (dashed line). Vertical red lines indicate the bounds set by HB stars for each individual model.

because of the large hierarchy in the SM quark masses. Without looking at explicit realisations of α, β and γ, it is neither possible to give exact bounds on the allowed F-masses nor to find the exact cutoff scale appearing the loop calculation. In conclusion, it is reasonable to expect $\Lambda \lesssim 10^3 f_a$, but larger values cannot be ruled out with absolute certainty.

Current bounds and sensitivity prospects of NA62 on the new models are plotted in figure 6.6. The $K^+ \to \pi^+ + a$ branching ratio as a function of f_a and for Λ in the range of $30f_a < \Lambda < M_{\text{Pl}}$ is shown as an orange band. The lower boundary of $\Lambda = 30f_a$ can be used to derive robust exclusions, since even smaller values of Λ and thereby m_F would lead to significant flavour effects already at tree level (see appendix C). The existing limits on the branching ratio from the E787 and E949 experiments thus rule out $f_a < 1.6 \times 10^6\,\text{GeV}$ in the new model. The future runs of NA62 could elevate this bound to $5.6 \times 10^6\,\text{GeV}$. Considering the possibility that Λ may be much larger than f_a leads to detection opportunities way above $f_a = 10^7\,\text{GeV}$. Only in this way, NA62 can go significantly beyond the robust limits set by the cooling of HB stars. This is illustrated by the red exclusion lines in figure 6.6. As discussed in the previous section, these bounds are much more restrictive for the EFT-ℓ model, which means that this model is nearly undetectable by NA62. As a comparison, the DFSZ predictions of the branching ratio is also included in the plot. It becomes clear that without the large logarithmic enhancement, there is no parameter space that can be tested by NA62 which is not already robustly excluded by astrophysical observations.

6.6 Conclusions

The starting point of the work presented in this chapter is a rather simple observation. When investigating flavour-dependent observables at one-loop level in an EFT framework, one requires an a priori undefined UV cutoff to regulate the divergence appearing in the EFT calculation. It is tempting and even common practice to infer this cutoff from the suppression scale appearing with every higher-dimensional operator in the EFT Lagrangian. This *leading log* prescription assumes that the EFT is valid up to the PQ symmetry breaking scale of order f_a. It is however in conflict with typical UV-complete QCD axion models.

By far the most prevalent models are the DFSZ and KSVZ models. The former does not reproduce the naive EFT result because it includes new degrees of freedom far below the PQ scale. These regulate the divergence in a similar fashion as new degrees of freedom in a Pauli-Villars regularization scheme. The DFSZ model therefore provides a physical interpretation of the cutoff, the mass of the additional charged Higgs particles $m_{H^\pm}$, but it is not in agreement with the large leading log. Hadronic axion models like the KSVZ, on the other hand, do not introduce new particles lighter than f_a since the masses of new fermions are generated by the spontaneous breaking of the PQ symmetry. In their case, it is the lack of tree-level axion interactions with SM fermions, which does not allow for a large logarithmically enhanced flavour effects.

The existence of this discrepancy demonstrates that in general the matching between observations and UV-complete theories is far from trivial. This applies in particular to the QCD axion because there is a large hierarchy of energy scales between potential axion observations and UV-complete models. A variety of new physics contributions could become important within this large energy range and result in a dependence of observations on parameters of the UV model. Furthermore, the question arises if a QCD axion model exists whose EFT description, containing only the axion as a new field, is valid up to f_a. In the previous sections, such a model was successfully constructed by combining electroweak and PQ spontaneous symmetry breaking to generate masses of all fermionic particles in the theory. This allows the SM particles to transform non-trivially under $\mathrm{U}(1)_{\mathrm{PQ}}$. Additional degrees of freedom appear in the form of fermionic messenger fields. These can obtain a large bare mass term without violating any gauge or global symmetries and can hence be assumed heavier than f_a. The new model therefore fulfils all requirements to reproduce EFT results.

Possible detection opportunities of the new model were discussed in section 6.5. The most relevant flavour observable is the $K^+ \to \pi^+ + a$ decay, which will be probed with unprecedented accuracy by the NA62 experiment. This is plotted in figure 6.6, which also shows that this detection technique is particularly relevant for the model version without axion lepton coupling. In this variant, the sensitivity to the rare kaon decay could supersede all astrophysical bounds except for the one based on observations of supernova SN1987A. But taking into account the uncertainties related to these limits[42], it becomes apparent that it would be highly desirable to

---

[42] An alternative scenario for the SN1987A explosion was considered in ref. [343], which would leave the SN bound on axions invalid. However, recent hints towards the existence of a neutron

confirm this bound by an entirely lab-based experiment. Another axion search, which could be sensitive to the version of the new model with tree-level lepton couplings, is the upcoming helioscope IAXO. It could reach close to $10^9\,\text{GeV}$ in f_a in its upgraded version IAXO+. This would supersede all astrophysical limits, including the previously mentioned supernova bound as well as the one derived from observations of the brightness of the tip of the red-giant branch in globular clusters.

The DM phenomenology of the investigated models is similar to those in the DFSZ and KSVZ benchmark models. It is important to note in this context that the models investigated above have a domain wall number, which is greater than one. This means that the scenario in which the PQ symmetry is broken during inflation and not restored afterwards is preferred. Alternatively, the discrete Z_N symmetry could be broken explicitly. While this can be useful to allow for the correct DM abundance at smaller values of the PQ scale f_a [292], potentially accessible with experiments such as NA62, it may also lead to a certain degree of tuning.

Despite mapping onto a rather minimal EFT – the property which was intended from the start – the new UV model has a rather complex flavour structure and in some parts of parameter space these could even become detectable. More detailed discussions on this topic as well as on potential CP-violating axion interactions can be found in appendix C. These types of couplings could also be employed in future studies to establish connections between the axion dynamics and the flavour puzzle. Furthermore, the new model is not uniquely defined by its EFT. In the course of its construction, several explicit choices had to be made. For example, it was necessary to fix the charges of the fields involved. Following alternative routes leads to similar models with some distinct differences. These are discussed in appendix B.

Coming back to the main topic of this thesis, the study presented in this chapter started by highlighting difficulties in connecting observations to axion models. It was made apparent that in order to make links between the two, it is necessary to develop a detailed understanding of the strength and limitations of EFT tools. This can be seen from the fact that it already turned out to be far from trivial to construct a UV embedding of a rather simple axion EFT that agrees with the leading log prescription. Inserting f_a as a cutoff scale by hand therefore spoils the general applicability of the calculation, which is supposed to be the main feature of EFTs. Nevertheless, a UV completion which agrees with the large cutoff was successfully constructed. This makes it possible to put stringent bounds on an explicit UV-complete axion model with a large logarithmic enhancement in flavour-violating decays. Even though a completion (with some theoretically disfavoured features) was found in this case, it remains an intriguing question whether extra assumptions – like the requirement of no additional fine-tuning compared to the SM – may severely restrict the space of axion EFTs that are the low-energy manifestations of reasonable UV models.

Certainly, the EFT approach is a cornerstone of axion phenomenology and will remain the default option for analysing new experimental data. However, the work outlined in this chapter demonstrates that it is highly beneficial to compare EFT results to common QCD axion models, which serve as benchmarks. This avoids

---

star remnant in the site of the explosion favour the standard hypothesis [344, 345].

misconceptions about the actual sensitivity of axion searches. In the case of the kaon decay, one would for instance note that the sensitivity in c_q/f_a differs by more than two orders of magnitude between the DFSZ and the EFT-like model. The new model could therefore be very useful as such a benchmark owing to its high UV cutoff and simple low-energy interactions. It thus provides another helpful link between experimental observations and UV-complete QCD axion models living at a much larger energy scale.

7 The solar axion spectrum

This chapter is based on results and arguments which were published in refs. [3] and [4] before submission of this thesis. Material and figures from these articles are used in the following without further reference.

The solar axion spectrum contains abundant information on the nature of the axion [195, 196, 233] as well as on the Sun itself [346, 347]. For this reason, this chapter is focused on the calculation of the spectral flux and detection opportunities of future solar axion searches.

Connections between EFTs, which are completely defined by a set of low-energy parameters, and UV-complete axion models were investigated in the previous two chapters. While it became apparent that this matching is far from trivial, the simple fact remains that good knowledge of as many as possible EFT parameters is crucial for successfully ruling out or favouring certain UV models. Solar axion searches are particularly well suited for this purpose. On the one hand, they are independent from cosmology in the sense that they do not rely on the existence, abundance or structure of axion dark matter. On the other hand, they are unique because of their large number of potentially detectable axion parameters including several couplings and the axion mass. This is due to various interactions with SM particles inside the Sun contributing to the total axion flux. The different components can be distinguished by their energy dependence, i.e. the resulting spectrum. On the detection side, the conversion probability of an axion into a photon depends on the energy and mass of the incoming axion. This allows to infer the mass on top of the coupling constants [195, 196].

In order to exploit the full potential of solar axion searches, it is crucial to precisely predict the expected spectral axion flux from the Sun. In addition, the remaining uncertainties should be quantified. This serves two purposes, which are central to axion phenomenology. In case of detection of a first signal, a good agreement between the measured and predicted spectrum would be an important consistency check in support of the axion interpretation of such a signal. Secondly, axion parameters can only be measured with a precision that is limited by the experimental uncertainty as well as the uncertainty of the theory prediction. In the following, the calculation of the solar axion flux is updated to include state-of-the-art solar models, plasma opacity data and all of the known and calculable corrections to the axion flux. For the first time, the uncertainty of such calculations is systematically quantified by estimating the known systematic effects related to approximations or known higher-order corrections and by propagating the statistical uncertainties of solar models and opacity data to the axion spectrum via a Monte Carlo simulation.

Additional motivation for the detailed investigation of the solar axion flux comes from the upcoming international axion observatory (IAXO) [193, 210]. It is expected to improve on the sensitivity of previous helioscopes [201–203, 205–207, 348–352] by more than an order of magnitude in the coupling constants. At the same time, it was shown that it could be sensitive enough to study solar properties like the abundances of metals [346] or magnetic fields inside the Sun [347]. However, both of these proposals hinge on the ability to accurately predict the solar axion flux. For the systematic investigation of uncertainties and the flexible calculation of the axion spectrum from solar models, a C++ code was developed which provides the required functionality.[43] All spectra and numerical results in this chapter were calculated with this code, which is also planned to be part of the signal and response calculation of IAXO.

7.1 Calculation of the solar axion flux

To review the calculation of the solar axion flux, it is first necessary to identify the relevant axion parameters. In the solar plasma, only interactions with electrons, photons and nuclei can conceivably play an important role. The latter can be most generically quantified by axion interactions with neutrons and protons, which will later be related to couplings to specific nuclear transitions. An axion EFT containing all significant parameters for solar axion searches is therefore given by

$$\mathcal{L}_{\odot} = \frac{1}{2}(\partial_\mu a)^2 - \frac{1}{2}m_a^2 a^2 - \frac{g_{a\gamma\gamma}}{4} a\, F_{\mu\nu}\widetilde{F}^{\mu\nu} + \sum_{f=e,n,p} \frac{g_{af}}{2m_f}\, (\partial_\mu a)\, \bar{f}\gamma^\mu\gamma^5 f\,, \tag{7.1}$$

where the sum in the fermion interactions term runs over electron, neutrons and protons. For the purposes of this chapter, the common phenomenological couplings constants $g_{a\gamma\gamma}$ and g_{ae}, which are proportional to the inverse of the axion decay constant, were chosen instead of the dimensionless $c_{\gamma\gamma}$ and c_f. Furthermore, the derivative basis for fermionic couplings is a convenient choice. This way there are no sizeable contributions to the axion-photon coupling from triangle diagrams of charged fermions if the momentum transfer is small compared to the mass of the electron [114]. This limit can safely be taken for all of the interactions considered below.

All axion production processes can be categorized by the coupling constant involved. At leading order in $1/f_a$ this can only be one of the three options – $g_{a\gamma\gamma}$, g_{ae} or nuclear interactions quantified by g_{an} and g_{ap}. All processes which are explicitly computed in this thesis are depicted in figure 7.1. Some additional contributions will be considered as corrections and their size is only parametrically estimated in section 7.3.

If a solar axion is supposed to be detectable in a helioscope, it must be lighter than $\sim 1\,\mathrm{eV}$ (see section 4.2). This motivates to only consider highly relativistic axions in the calculation of axion production rates because they are produced with typical energies of $\omega \sim \mathrm{keV} \gg m_a$. All of the following calculations are performed in this limit.

---

[43] The code has been made available as open source software at `https://github.com/sebhoof/SolarAxionFlux`.

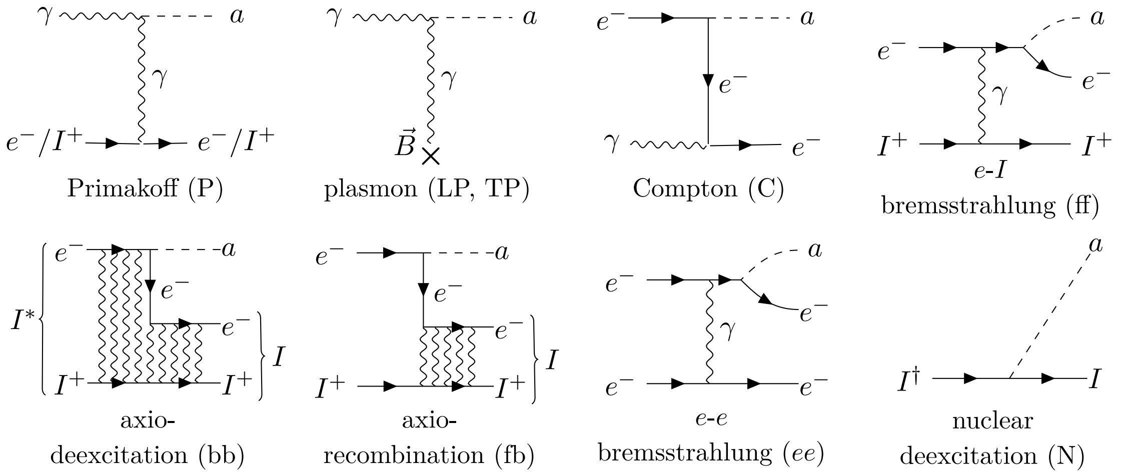

Figure 7.1. Diagrams of all relevant axion production processes [13, 105, 231] inside the Sun. In the second diagram, $\vec{B}$ denotes the macroscopic magnetic field inside the Sun. I stands for an arbitrary chemical element. It can be in an atomic (I^*) or nuclear ($I^\dagger$) excited state. Electromagnetically bound states are indicated by many photon lines. The two diagrams including this illustration are not to be strictly understood as Feynman diagrams. For simplicity, only simply-charged ions or neutral elements are depicted, but all processes can also occur with elements in higher ionisation states.

7.1.1 Primakoff effect in a non-degenerate plasma

The Primakoff effect is the most relevant production mechanism of axions coupled to photons. The Feynman diagram of the process is shown in the top left corner of figure 7.1. It can be understood from a semi-classical standpoint as an excitation of the electromagnetic field inside the plasma – called a plasmon – converting into an axion in the electromagnetic field of electrons or ions.

This effect has always been the focus of solar axion searches, because a non-vanishing tree-level coupling to photons is a generic prediction of QCD axions due to the model-independent contribution in equation (3.31). Furthermore, helioscope detection anyway requires a non-vanishing $g_{a\gamma\gamma}$ and therefore it is the only coupling constant which can be constrained individually.

The first calculation of the Primakoff production rate per axion phase space (Γ_a^{P}) to include the charge screening effect in the solar plasma was performed by Raffelt [231, 353]. His result can be expressed in terms of the Debye screening scale κ_{s}, the plasma temperature T, the axion energy ω, the number density of free electrons n_e and the axion-photon coupling $g_{a\gamma\gamma}$ as

$$\Gamma_a^{\mathrm{P}}(\omega) = g_{a\gamma\gamma}^2 \alpha_{\mathrm{EM}} \, \frac{n_e + \bar{n}}{8} \left[\left(1 + \frac{\kappa_{\mathrm{s}}^2}{4\omega^2} \right) \ln \left(1 + \frac{4\omega^2}{\kappa_{\mathrm{s}}^2} \right) - 1 \right] \frac{2}{e^{\omega/T} - 1} \, . \tag{7.2}$$

$\bar{n}$ is defined as the sum of number densities of each ion (n_z), weighted by the square of its electric charge (Q_z^2) in units of the elementary charge squared,

$$\bar{n} \equiv \sum_z Q_z^2 \, n_z \, . \tag{7.3}$$

The index z labels all different types of ions in the plasma. It is important to note that not all nuclei are fully ionised even in the centre of the Sun and therefore Q_z is not simply given by the atomic number. How to evaluate $\bar{n}$ is further discussed in section 7.1.6. In the non-degenerate limit, the Debye scale, which quantifies the size of the screening charge cloud, can be computed from the temperature and densities as [231]

$$\kappa_{\mathrm{s}}^2 = \frac{4\pi\alpha_{\mathrm{EM}}}{T}(n_e + \bar{n})\,. \tag{7.4}$$

Note that in order to arrive at equation (7.2), Raffelt worked in the limit of a static screening potential during each scattering event. A first-order approximation beyond this strict limit is introduced in appendix E. The resulting corrections are parametrically estimated and enter the list of systematic uncertainties in table 7.3.

The rate in equation (7.2) also assumes a linear dispersion relation for the scattering photons. This is however not the case inside a plasma, in which the dispersion relation in the non-relativistic[44] and non-degenerate limit is given by [233]

$$\omega^2 = \omega_{\mathrm{pl}}^2 + k_\gamma^2\,, \tag{7.5}$$

where k_γ is the photon momentum, which is identical to ω in vacuum, and ω_{pl} denotes the plasma frequency. This relation is also not exact, but higher-order corrections are suppressed by powers of temperature over electron mass T/m_e. The size of these corrections and their impact are discussed in section 7.3. ω_{pl} is given by [233]

$$\omega_{\mathrm{pl}}^2 = \frac{4\pi\alpha_{\mathrm{EM}}}{m_e}\,n_e\,, \tag{7.6}$$

The non-trivial dispersion relation of plasmons can have a large effect when $\omega \sim \omega_{\mathrm{pl}}$ and therefore equation (7.2) is only valid in the limit $\omega \gg \omega_{\mathrm{pl}}$. In fact, the rate has to vanish for $\omega < \omega_{\mathrm{pl}}$ since plasmons of energy smaller than the plasma frequency do not exist. The full result including the dispersion relation (7.5) was found to be [347, 354]

$$\Gamma_a^{\mathrm{P}}(\omega) = g_{a\gamma\gamma}^2\alpha_{\mathrm{EM}}\,\frac{n_e+\bar{n}}{8}\left[\frac{1}{2}\int_{-1}^{1}\mathrm{d}x\,\frac{1-x^2}{(\xi_1-x)(\xi_2-x)}\right]\frac{2}{e^{\omega/T}-1}\frac{\mathrm{d}k_\gamma}{\mathrm{d}\omega} \tag{7.7}$$

$$\text{with}\quad \xi_1 \equiv \frac{k_a^2+k_\gamma^2}{2k_ak_\gamma}\quad\text{and}\quad \xi_2 \equiv \xi_1 + \frac{\kappa_{\mathrm{s}}^2}{2k_ak_\gamma}\,. \tag{7.8}$$

k_a denotes the axion momentum, which is equal to ω in the relativistic limit. For the numerical evaluation, it is useful to perform the integral analytically. This results in the rate

$$\begin{aligned}\Gamma_a^{\mathrm{P}}(\omega) =& g_{a\gamma\gamma}^2\alpha_{\mathrm{EM}}\,\frac{n_e+\bar{n}}{8}\frac{2}{e^{\omega/T}-1}\frac{\omega}{k_\gamma}\\ &\times\left[\frac{(\xi_1^2-1)\ln\left(\frac{\xi_1+1}{\xi_1-1}\right)-(\xi_2^2-1)\ln\left(\frac{\xi_2+1}{\xi_2-1}\right)}{2\,(\xi_1-\xi_2)}-1\right].\end{aligned} \tag{7.9}$$

---

[44] *Non-relativistic* refers to the velocities of charges, i.e. electrons and nuclei in the solar plasma. The electromagnetic field excitations called plasmons can still propagate at highly relativistic speeds.

Note that this equation is only valid for $\omega > \omega_{\mathrm{pl}}$ and it goes to zero in the limit $\omega \to \omega_{\mathrm{pl}}$. For smaller ω, the rate can be directly set to zero.

All in all, the Primakoff production rate in a non-degenerate and non-relativistic plasma can be expressed in a closed form and, when all definitions introduced above are inserted, it only depends on the three solar quantities T, n_e and $\bar{n}$. The computation of the last two of these is discussed in section 7.1.6.

7.1.2 Primakoff effect including partial degeneracy

The calculation of the Primakoff rate in the previous section assumed free charges to be non-degenerate inside the Sun. The electron gas is however partially degenerate [233], which can have an impact at the percent level on quantities like the Debye scale, the plasma frequency as well as the total scattering rate.

In order to include these corrections, one has to evaluate the phase space integrals of initial and final state electrons using the full Fermi-Dirac distribution,

$$f(p) = \frac{1}{e^{(E(p)-\mu)/T} + 1} \, , \tag{7.10}$$

where E denotes the energy as a function of momentum p and μ is the chemical potential. From this, one can find the total number density of electrons in the non-relativistic limit as [355]

$$n_e = 2 \left(\frac{m_e T}{2\pi} \right)^{3/2} \mathcal{F}^{\mathrm{FD}}_{1/2}(z) \, , \tag{7.11}$$

where $z = (\mu - m_e)/T$ and the Fermi-Dirac integral is defined as

$$\mathcal{F}^{\mathrm{FD}}_{j}(z) \equiv \frac{1}{\Gamma(j+1)} \int_0^\infty \mathrm{d}t \, \frac{t^j}{e^{t-z} + 1} = - \sum_{k=1}^{\infty} (-1)^k \, \frac{\exp(z)^k}{k^{j+1}} \, . \tag{7.12}$$

Γ denotes the usual gamma function. The two parameters n_e and T, which are required for the computation of the Primakoff rate, are given by the solar model. This means that in order to find the distribution function $f(p)$, one has to solve equation (7.11) for μ, which is a straightforward numerical task.

Once $f(p)$ is known, the full expression for the screening scale including the degeneracy can be evaluated. It is given by [233]

$$\kappa_{\mathrm{s}}^2 = \frac{4\pi\alpha_{\mathrm{EM}}}{T} \bar{n} + \frac{4\alpha_{\mathrm{EM}}}{\pi} \int_0^\infty \mathrm{d}p \, f(p) \, p \left(v + v^{-1} \right) \, , \tag{7.13}$$

where $v = p/E = p/\sqrt{m_e^2 + p^2}$ is the electron's velocity. Note that the contribution from ions in the plasma is unaffected. This is because their mass is at least a factor of 2000 larger than m_e and they are thus highly non-degenerate. One can also check that in the Boltzmann limit of $f(p) \propto e^{-p^2/2m_e}$ the expression in equation (7.4) is reproduced. Degeneracy effects become most important close to the core of the Sun, where the largest reduction of κ_{s} compared to equation (7.4) is numerically found to be 1.2 %. The correction is comparably small because electrons only contribute ∼37 % to the total value of κ_{s}^2 in the core. This is illustrated in figure 7.2.

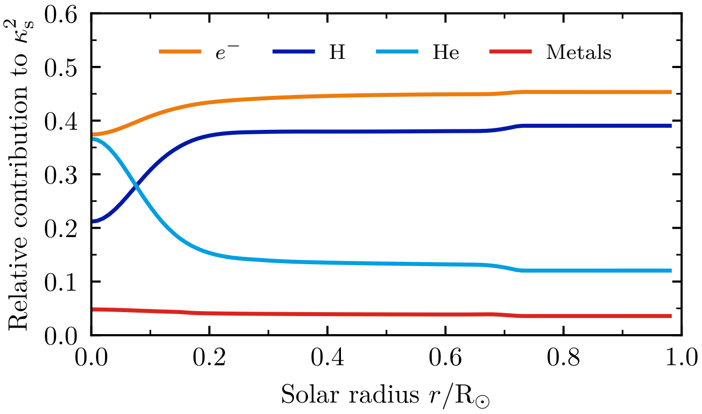

Figure 7.2. Relative contributions of electrons, H, He and heavier elements (metals) to κ_s^2 according to equation (7.13) for the AGSS09 model.

The same rationale as for the Debye scale can be applied to the plasma frequency. ω_{pl} is given in terms of $f(p)$ with an arbitrary degree of degeneracy as [233]

$$\omega_{\mathrm{pl}}^2 = \frac{4\alpha_{\mathrm{EM}}}{\pi} \int_0^\infty \mathrm{d}p\, f(p)\, p \left(v - \frac{1}{3} v^3 \right) . \tag{7.14}$$

The crucial difference compared to the Debye scale is that only positive powers of the velocity v appear in the integrand. Because degeneracy effects are strongest at small velocities, i.e. for low-energy states, the plasma frequency is less affected by the partial degeneracy even though it does not receive a contribution from the non-degenerate ions. The largest relative difference between the full expression and the approximation in equation (7.6) is found to be of order 10^{-3}. Nonetheless, all of the following numerical calculations work with the full expression including degeneracy corrections.

Another effect of the degeneracy is the occupation of phase space for outgoing particles in scattering processes – called Pauli blocking. The electron occupation numbers can reach up to ∼0.2 in the solar core, which makes this suppression non-negligible for processes including final states with small momentum. To compute the size of the effect, one can apply a similar strategy like the one in ref. [233] in the context of Coulomb scattering. To start with, a suppression factor F_{deg} is defined, which reduces the rate of scattering events with electrons by replacing

$$n_e \mapsto F_{\mathrm{deg}}\, n_e , \tag{7.15}$$

in equation (7.9). The suppression via Pauli blocking is given by the ratio of scattering rates with and without a factor of $(1 - f(p_2))$ in the phase space integrals,

$$F_{\mathrm{deg}} = \frac{\int_0^\infty \mathrm{d}p_1 \int_{-1}^1 \mathrm{d}y \int_{-1}^1 \mathrm{d}x\; f(p_1)\; (1 - f(p_2))\, \frac{\mathrm{d}\sigma}{\mathrm{d}\Omega}}{\int_0^\infty \mathrm{d}p_1 \int_{-1}^1 \mathrm{d}y \int_{-1}^1 \mathrm{d}x\; f(p_1)\, \frac{\mathrm{d}\sigma}{\mathrm{d}\Omega}} , \tag{7.16}$$

where p_1 and p_2 are the incoming and outgoing electron momenta, respectively. The integrals over $y \equiv \cos\theta_{12}$ and $x \equiv \cos\theta$ are averages of the angle between the incoming momentum and the momentum transfer θ_{12}, as well as the scattering

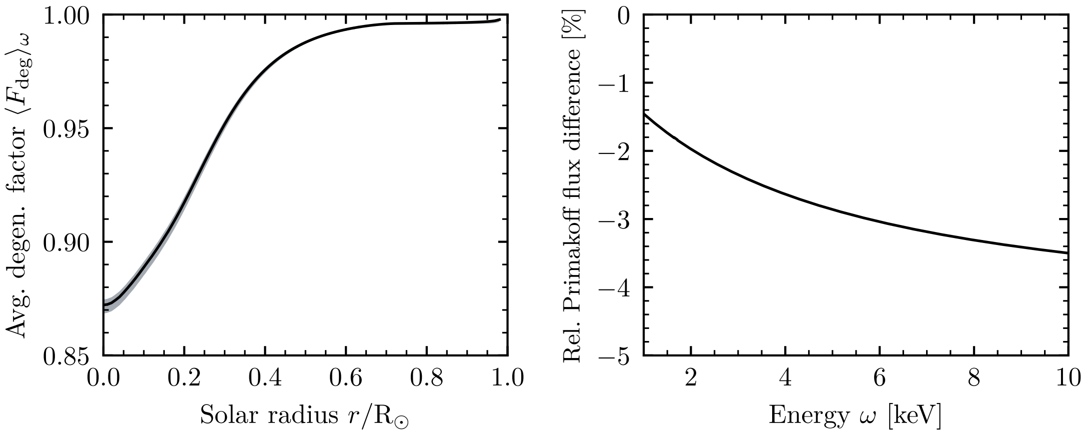

Figure 7.3. The effect of partial electron degeneracy on the Primakoff production rate *Left:* Energy-averaged electron degeneracy factor for the solar model AGSS09 (black line). The shaded grey region indicates the range of values of F_deg between 1 keV–10 keV. *Right:* Primakoff flux reduction (in percent) due to *all* degeneracy effects in the energy range of interest and for the AGSS09 model.

angle θ of the plasmon converting to an axion. $f(p)$ is the Fermi-Dirac distribution as in equation (7.10) and $\frac{\mathrm{d}\sigma}{\mathrm{d}\Omega}$ is the differential Primakoff cross section from the integrand of equation (7.7),

$$\frac{\mathrm{d}\sigma}{\mathrm{d}\Omega} = \frac{\alpha_\mathrm{EM}}{32\pi}\, g_{a\gamma\gamma}^2 \, \frac{1-x^2}{(\xi_1 - x)(\xi_2 - x)} . \tag{7.17}$$

Energy and momentum conservation lead to the relations,

$$p_2^2 = p_1^2 + 2\, p_1 q \cos\theta_{12} + q^2 \, , \tag{7.18}$$

and

$$\left(\frac{q}{\omega}\right)^2 = 2 - \left(\frac{\omega_\mathrm{pl}}{\omega}\right)^2 - 2\sqrt{1 - \left(\frac{\omega_\mathrm{pl}}{\omega}\right)^2}\cos\theta \, , \tag{7.19}$$

where q is the transferred momentum and ω the axion energy. By plugging these identities into the expression for F_deg, the integrals can be computed numerically[45] for a given distance from the solar core r and axion energy ω. At the centre of the Sun, F_deg is found to be ~ 0.87 with only weak dependence on ω in the relevant energy range of 1 keV to 10 keV. To slightly simplify the calculation, it is sensible to average over F_deg in this energy range and to assume it to be a constant in ω from now on. The range of true values as well as the average as a function of r is given in the left panel of figure 7.3.

Finally, all of the degeneracy effects mentioned in this section can be combined to find the relative reduction of the total Primakoff flux as a function of axion energy. This is plotted in the right panel of figure 7.3. The integrated Primakoff flux in the relevant energy range is reduced by about 2.6 %. These effects have previously not been accounted for.

---

[45] Note that the integrals over $x = \cos\theta$ in both the numerator and denominator, as well as the integral over $y = \cos\theta_{12}$ in the denominator, can be performed analytically to reduce the dimensionality of the numerical integral.

7.1.3 Plasmon conversion in the solar magnetic field

In addition to the electromagnetic interactions on atomic distance scales, the Sun is also permeated by a large scale macroscopic magnetic field. Inside of this field, propagating plasmons can convert to axions, which is the reverse of the process taking place in axion helioscopes. This production mechanism has lately gained renewed attention because axions from plasmon conversion may come within reach of observations in the not-so-distant future [347, 356] and because it may be relevant for stellar energy loss arguments [356, 357].

Because of the effective mass of plasmons (see the non-linear dispersion relation in equation (7.5)), they exist in both transverse and longitudinal polarisations, which have to be considered separately. The conversion rate of longitudinal plasmons into axions was found to be [347, 356]

$$\Gamma_a^{\mathrm{LP}}(\omega) = \frac{g_{a\gamma\gamma}^2 B_\parallel^2}{e^{\omega/T} - 1} \frac{\omega^2 \Gamma_{\mathrm{L}}}{(\omega^2 - \omega_{\mathrm{pl}}^2)^2 + (\omega \Gamma_{\mathrm{L}})^2} , \tag{7.20}$$

where $B_\parallel$ is the magnetic field strength projected onto the propagation direction of the plasmon and Γ_{L} is the total collision rate of longitudinal plasmons with any scattering partner in the plasma. The resulting rate Γ_a^{LP} is dominated by the resonance at $\omega = \omega_{\mathrm{pl}}$ and Γ_{L} merely defines the width of this peak, which is typically too narrow to be resolved. Because the total flux is independent of Γ_{L} a detailed evaluation of its value is not strictly necessary.[46] For the same reason, Γ_a^{LP} was approximated as a delta function in refs. [347, 356, 357]. Since the energy of axions produced in this process is defined by the plasma frequency, which reaches a maximum of $\omega_{\mathrm{pl}} \lesssim 0.3\,\mathrm{keV}$ in the solar core, longitudinal plasmon conversion only contributes at small energies. The resulting spectrum is not monochromatic but features peaks. Their position is given by the plasma frequency at radii with the strongest magnetic field.

Far from resonance, the conversion rate of longitudinal plasmons is strongly suppressed. This is due to their peculiar dispersion relation. In the limit of non-relativistic, non-degenerate electrons and to leading order in T/m_e, it is [233]

$$\omega^2 = \omega_{\mathrm{pl}}^2 \left(1 + 3\,\frac{k^2}{\omega^2}\,\frac{T}{m_e} \right) . \tag{7.21}$$

Because $T \ll m_e$, the second term in parentheses can be neglected for momenta of the order of the energy and the relation reduces to $\omega \approx \omega_{\mathrm{pl}}$ [233, 357]. This is exactly the resonance condition. The second term only becomes relevant, when the phase velocity k/ω reaches values of the order of the thermal velocity of electrons $v^* \sim \sqrt{T/m_e}$ [233]. At this point the plasmon four momentum crosses the *electron cone*, which means that a large number of electrons fulfils the Cherenkov condition. This causes Landau damping [233, 358] and leads Raffelt to the conclusion that no organised oscillations exist at $k \gtrsim \omega/v^*$ [233]. For the calculation of the axion

---

[46] The tiny width of the resonance can lead to very slow convergence of numerical integrals. Γ_{L} was therefore set to $\max(\Gamma_{\mathrm{T}}, 0.1\,\mathrm{eV})$ by hand, which improves convergence and does not alter the total number of axions produced.

flux, this means that it is sufficient to integrate equation (7.20) over a small region around the resonance. All additional off-resonance contributions are expected to be completely negligible.

The behaviour of transverse plasmons is much closer to the usual intuition. Their dispersion relation is given in equation (7.5), which illustrates that the plasma frequency acts as an effective mass of transverse electromagnetic excitations in the plasma. As a consequence, resonant conversion is only possible when the axion is massive with $m_a = \omega_{\mathrm{pl}}$. The corresponding rate is given by [347]

$$\Gamma_a^{\mathrm{TP}}(\omega) = \frac{2\,g_{a\gamma\gamma}^2 B_\perp^2}{e^{\omega/T} - 1} \frac{\omega^2 \Gamma_{\mathrm{T}}}{(\omega_{\mathrm{pl}}^2 - m_a^2)^2 + (\omega \Gamma_{\mathrm{T}})^2} \,. \tag{7.22}$$

$B_\perp^2$ is the average square of the magnetic field projected onto the polarisation vector as defined in ref. [347]. Γ_{T} is the collision rate of transverse plasmons and the prefactor of two accounts for the two linearly independent polarisation states. Because this chapter is only concerned with axions which may be detectable in helioscopes, the possible resonance at $m_a = \omega_{\mathrm{pl}} \sim 100\,\mathrm{eV}$ is of no interest. Instead, equation (7.22) is taken in the limit $m_a \to 0$, where the axion production rate becomes approximately proportional to the collision rate Γ_{T}. This means that unlike for the longitudinal plasmons, the rate Γ_{T} is highly relevant for computing the off-resonance contribution of transverse plasmons. In addition, Γ_{T} is a frequency-dependent quantity. It can, however, be deduced from the monochromatic opacity $\kappa(\omega)$, which is defined as the absorption coefficient k (inverse of the mean free path) per plasma density ρ,

$$\kappa(\omega) = \frac{k(\omega)}{\rho} \,. \tag{7.23}$$

The total collision rate is then given by [357]

$$\Gamma_{\mathrm{T}} = k(\omega) \left(1 - e^{-\omega/T}\right) = \kappa(\omega)\,\rho \left(1 - e^{-\omega/T}\right) \,. \tag{7.24}$$

Details about different opacity codes are given in sections 7.1.4 and 7.2.

The calculation in this thesis, which was published before submission in ref. [3], is the first to incorporate monochromatic opacities. Previously, the Rosseland mean opacity κ_{R} has been used for all frequencies [357]. It is defined by an integral over the monochromatic opacity as [359]

$$\frac{1}{\kappa_{\mathrm{R}}} = \int_0^\infty \frac{\mathcal{R}(\omega/T)}{T\,\kappa(\omega)}\,\mathrm{d}\omega \,, \tag{7.25}$$

where $\mathcal{R}$ is the Rosseland weight function,

$$\mathcal{R}(u) \equiv \frac{15}{4\pi^4} \frac{u^4\,e^u}{(e^u - 1)^2} \,. \tag{7.26}$$

Rosseland opacities are a powerful tool for computing the radiative energy transport in stellar plasmas. For the purpose of calculating the off-resonance spectral axion flux, it is however not justified to replace the monochromatic $\kappa(\omega)$ by a frequency averaged quantity. This distinction is highly relevant because the two results differ by

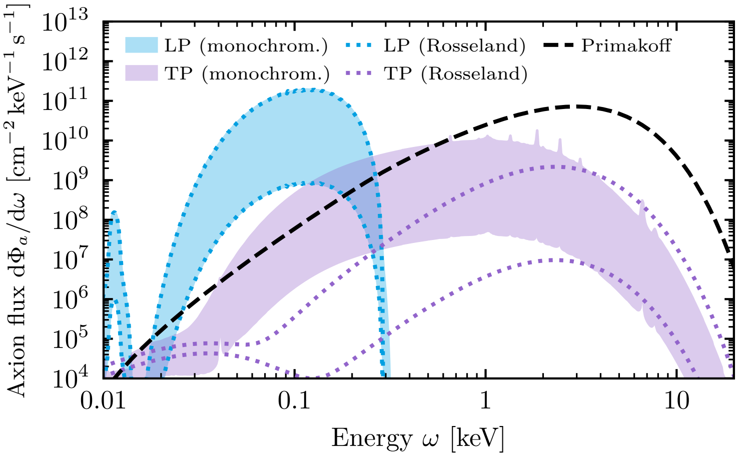

Figure 7.4. Axion flux from longitudinal (LP) and transverse (TP) plasmon interactions for $g_{a\gamma\gamma} = 10^{-10}\,\mathrm{GeV}^{-1}$. The blue and purple shaded contours show the LP and TP flux between the minimal and maximal B-field reference values, respectively, using monochromatic opacities $\kappa(\omega, r)$. Dotted lines show the LP and TP fluxes computed with Rosseland mean opacities. The Primakoff flux is plotted as a dashed black line for comparison.

more than three orders of magnitude at specific energies. In addition, the monochromatic opacity has some distinct features like peaks from atomic transitions. These should also be visible in the resulting axion flux as shown in figure 7.4.

To finally compute the spectral flux from plasmon conversion requires a model of the internal solar magnetic field. Despite the fact that many solar quantities are known with high precision (see section 7.2), the magnetic field is still poorly constrained [347, 360]. This is also why it has been suggested to use the axion flux as a probe of solar magnetic fields instead of the other way around [347]. In this work, the magnetic field model used in ref. [347] is adopted for concreteness and in order to obtain comparable results. It postulates a toroidal field with

$$\vec{B}(\vec{r}) = -3\,B(r)\cos\theta\,\sin\theta\;\hat{e}_\phi\;, \tag{7.27}$$

where θ is the azimuthal angle inside the Sun. The normalization is given by

$$B(r) = \begin{cases} B_\text{rad}(1+\lambda)(1+\frac{1}{\lambda})^\lambda \left(\frac{r}{r_\text{CZ}}\right)^2 \left[1-\left(\frac{r}{r_\text{CZ}}\right)^2\right]^\lambda & \text{for } r < r_\text{CZ} - d_\text{tach} \\ B_\text{tach}\left[1-\left(\frac{r-r_\text{CZ}}{d_\text{tach}}\right)^2\right] & \text{for } |r - r_\text{CZ}| < d_\text{tach} \\ B_\text{outer}\left[1-\left(\frac{r-r_\text{upper}}{d_\text{upper}}\right)^2\right] & \text{for } |r - r_\text{upper}| < d_\text{upper} \\ 0 & \text{otherwise} \end{cases}, \tag{7.28}$$

$$\text{with} \qquad \lambda \equiv 10\,\frac{r_\text{CZ}}{\mathrm{R}_\odot} + 1\;. \tag{7.29}$$

Here, $r_\text{CZ} \approx 0.712\,\mathrm{R}_\odot$ is the approximate radius of the radiative zone, $r_\text{upper} \approx 0.96\,\mathrm{R}_\odot$ is the nominal beginning of the outer layers of the Sun while $d_\text{tach} \approx 0.02\,\mathrm{R}_\odot$ and $d_\text{upper} \approx 0.035\,\mathrm{R}_\odot$ are the shell thickness of the tachocline and the outer layers of the Sun, respectively. Again, for comparability, the same ranges for the B-field normalizations as considered in figure 3 of ref. [347] are chosen. These are $B_{\text{rad}} \in [200\,\text{T},\, 3000\,\text{T}]$, $B_{\text{tach}} \in [4\,\text{T},\, 50\,\text{T}]$ and $B_{\text{outer}} \in [3\,\text{T},\, 4\,\text{T}]$.

The results for the spectral axion flux from conversions of longitudinal and transverse plasmons[47] are plotted in figure 7.4. For comparison, the Primakoff flux and the approximation using Rosseland opacities are also shown. The large possible range of plasmon flux values for a given energy is due to the fact that the normalization of the large-scale solar magnetic field is so poorly constrained.

Like the Primakoff effect, plasmon conversions are due to $g_{a\gamma\gamma}$ and the resulting flux is subject to large uncertainties. They should, however, not simply be regarded as corrections or sources of uncertainty of the Primakoff prediction. This is because the peaks in the axion spectrum corresponding to longitudinal plasmon conversion are well separated from the bulk of the Primakoff contribution and they can therefore be detected independently. As it was pointed out in ref. [347], this would enable a measurement of the magnetic fields in the deep solar interior. Such a measurement would also allow for a more accurate prediction of the non-resonant conversion rate of transverse plasmons. In addition, the angular distribution of the flux from plasma conversions depends on the geometry of the solar magnetic field. This results in an anisotropic axion emission and an annually oscillating flux on Earth, providing another handle to distinguish plasmon conversions from the Primakoff flux [347]. In the following, the uncertainty of the Primakoff flux is therefore evaluated without taking plasmon conversion into account. Nonetheless, the relative strength of the two fluxes at specific energies and averaged over the whole spectrum is given in table 7.3.

7.1.4 Axion-electron interactions

The solar plasma contains large numbers of free and bound electrons constantly interacting. It is therefore natural to expect various relevant production processes of axions if they are coupled to electrons. In fact, all SM scattering processes taking place inside the Sun have an axion equivalent. This is the same process but with an outgoing photon or plasmon being replaced by an axion. Because the axion is a scalar particle, the spin structure of these processes differs from their SM counterpart. The diagrams of the relevant contributions involving axion-electron interactions are shown in figure 7.1. In combination, the resulting flux is often referred to as ABC flux.

Some of the production rates, such as electrons scattering off nuclei (ff), electron bremsstrahlung (ee) and the Compton effect (C) can be computed analytically similar to the Primakoff rate. As before, screening of scattering potentials has to

---

[47] The result in figure 7.4 differs from the one in ref. [347] by a constant factor. This due the geometrical factor, which is required to calculate the time averaged flux on Earth from Γ_a^{LP} as discussed in section III D of ref. [347]. It is left out here because it depends on the data taking times throughout the year and therefore is – at least slightly – different for each experiment. Instead, the flux averaged over all emission directions from the Sun is shown, for which the geometric factor was found to be 1/3 in ref. [357]. To recover the result in ref. [347], one has to multiply the spectrum labelled LP in figure 7.4 by 3×1.8.

be taken into account for processes taking place inside a plasma. The most recent evaluations of these rates were given in ref. [105] as

$$\Gamma^{\text{ff}}_{a,z}(\omega) = g_{ae}^2\,\alpha_{\text{EM}}^2\,\frac{8\pi}{3\sqrt{2\pi}}\,\frac{Q_z^2 n_z n_e}{\sqrt{T} m_e^{7/2}\omega}\,e^{-\omega/T}\,\mathcal{I}(\omega/T, y)\,, \tag{7.30}$$

$$\Gamma^{\text{ff}}_{a}(\omega) = \sum_z \Gamma^{\text{ff}}_{a,z}\,, \tag{7.31}$$

$$\Gamma^{ee}_{a}(\omega) = g_{ae}^2\,\alpha_{\text{EM}}^2\,\frac{4\sqrt{\pi}}{3}\,\frac{n_e^2}{\sqrt{T} m_e^{7/2}\omega}\,e^{-\omega/T}\,\mathcal{I}(\omega/T, \sqrt{2}y)\,, \tag{7.32}$$

$$\Gamma^{\text{C}}_{a}(\omega) = g_{ae}^2\,\alpha_{\text{EM}}\,\frac{n_e}{3m_e^4}\,\frac{\omega^2}{e^{\omega/T}-1}\,, \tag{7.33}$$

where the auxiliary screening function is defined as[48]

$$\mathcal{I}(u, y) = \int_0^\infty \mathrm{d}x\, x\, e^{-x^2} \int_{\sqrt{x^2+u}-x}^{\sqrt{x^2+u}+x} \mathrm{d}t\, \frac{t}{(t^2+y^2)}\,, \tag{7.34}$$

and the parameter y is given by

$$y \equiv \frac{\kappa_{\text{S}}}{\sqrt{2m_e T}}\,. \tag{7.35}$$

z runs again over all types of nuclei. The only difference to ref. [105] is that this thesis uses a different screening prescription, which changes the integrand in the function $\mathcal{I}$. This is because charge screening for bremsstrahlung is best described by an effective form factor $q^2/(q^2+\kappa_{\text{S}}^2)$ in the transition amplitude. q is the transferred momentum, i.e. the momentum of the exchanged plasmon. Ref. [105] used the square of this factor. The two choices correspond to either taking the average over charge distributions at the level of the matrix element squared or the matrix element itself. Following the arguments given in refs. [231, 233, 361], the former description is better if the time it takes for the scattering electron to cross the potential is smaller than the typical timescale of screening cloud formation. While both are approximations, the one used in this thesis should be more appropriate for bremsstrahlung processes. This is taking into account that slow moving ions are responsible for the bulk of the screening effect. Nevertheless, a more systematic investigation of the form factor could be beneficial. For the purpose of this thesis, it suffices to numerically compare the two screening prescriptions. The integrated ABC flux following equations (7.30) to (7.33) between 1 keV and 10 keV is about 5.7 % larger compared to the one with the form factor used in ref. [105]. The spectral flux within the same energy range differs by at most 26 %.

The partial degeneracy of electrons can also affect the rates above. On the one hand, the screening scale is modified as described in section 7.1.2. This effect is included in all of the following results. On the other hand, screening leads to significant Pauli blocking because all three processes contain at least one electron in the final state (see also refs. [233, 362]). The size of the resulting correction is estimated in section 7.3 and contributes to the overall uncertainty of the ABC flux.

---

[48] The second integral can be computed analytically, which is beneficial for numerical evaluations of the production rates.

All of the processes mentioned so far do not depend on the details of the atomic structure, which makes the calculation of these production rates comparably simple. This is not the case for bound-bound (bb) and free-bound (fb) transitions. A first spectral evaluation of the corresponding rates was performed by Redondo in ref. [105] by relating the axion production to the monochromatic opacity κ as defined in equation (7.23). This can be done in two steps. First, the photon absorption rate (quantified by the opacity) is related to the photon production by detailed balance under the assumption of local thermal equilibrium [105]

$$\Gamma^{i}_{\gamma,\,\mathrm{abs}}\, e^{-\omega/T} = \Gamma^{i}_{\gamma,\,\mathrm{prod}} \, . \tag{7.36}$$

$\Gamma^{i}_{\gamma,\,\mathrm{abs}}$ and $\Gamma^{i}_{\gamma,\,\mathrm{prod}}$ are the absorption and production rates per photon phase space for process i (with $i = \mathrm{bb}, \mathrm{fb}, \mathrm{ff}$), respectively. The factor $e^{-\omega/T}$ is required in order to have the same overall production and absorption rates per solar volume, taking into account stimulated emission as well as the thermal occupation number. All of the following rates will be production rates, which is why the subscripts are dropped. In a second step, the ratio of axion to spin averaged photon production was found for each of the three types of atomic processes as [105]

$$\frac{\Gamma^{i}_{a}}{\Gamma^{i}_{\gamma}} = \frac{1}{8\pi}\frac{g_{ae}^{2}\omega^{2}}{\alpha_{\mathrm{EM}} m_e^{2}} \, . \tag{7.37}$$

So by going from absorption to production and from photons to axions, the total ABC emission rate of axions Γ_a^{ABC} can be expressed in terms of the monochromatic absorption coefficient $k(\omega)$ as well as the Compton and ee bremsstrahlung rates [105],

$$\begin{aligned} \Gamma_a^{\mathrm{ABC}} &\equiv \Gamma_a^{\mathrm{ff}} + \Gamma_a^{\mathrm{fb}} + \Gamma_a^{\mathrm{bb}} + \Gamma_a^{\mathrm{C}} + \Gamma_a^{ee} & (7.38)\\ &= \frac{1}{8\pi}\frac{g_{ae}^{2}\omega^{2}}{\alpha_{\mathrm{EM}} m_e^{2}}\frac{k(\omega)}{e^{\omega/T}-1} + \frac{1}{2}\frac{e^{\omega/T}-2}{e^{\omega/T}-1}\Gamma_a^{\mathrm{C}} + \Gamma_a^{ee} \, . & (7.39) \end{aligned}$$

For details of the derivations the reader is referred to ref. [105]. This final expression can be used whenever the total monochromatic opacity of a plasma is available. Luckily, the opacity is a crucial ingredient for solar modelling since it determines the amount of energy transport by radiative diffusion. Accordingly, large efforts have been made to compute opacities of stellar plasmas (see refs. [359, 363–367] for an overview of works in this direction). These calculations include the aforementioned atomic structures as well as plasma effects. Opacity tables calculated by the opacity project (OP) even quote the contributions of each element separately. In this case, Redondo suggests [105] to calculate the ff processes involving hydrogen and helium directly with equation (7.31) and only include the opacities from elements heavier than helium in the first term of equation (7.39) while adding the full Compton rate. This thesis follows this suggestion whenever possible.

7.1.5 Axion-nucleon interactions

The final axion interactions considered for the calculation of the solar axion flux are the ones to nucleons. In principle, these allow two types of production mechanisms to take place. Axion emission in nuclear fusion or decay processes as well as thermal excitation and subsequent deexcitation of nuclei of stable isotopes.

Axions produced in nuclear fusion and decay processes typically carry an energy of the order of MeV. For instance, the axion flux from the $p+d \to {}^{3}\mathrm{He}+a\,(5.5\,\mathrm{MeV})$ reaction, which provides one of the most intense axion fluxes from nuclear reactions, has been experimentally searched for using Borexino [368] and SNO [369] data. In principle, helioscopes could also detect such high-energetic axions when they are equipped with a gamma-ray detector [204], but the resulting bounds are relatively weak.

The second option turns out to be more promising. Some isotopes with a significant abundance inside the Sun feature low-lying nuclear transitions, which can be thermally excited. In the past, the two candidates ${}^{57}\mathrm{Fe}$ [370, 371] and ${}^{83}\mathrm{Kr}$ [372] were identified. The corresponding transition energies E^* are 14.4 keV and 9.4 keV, respectively. In comparison to the solar core temperature of $\sim 1.3\,\mathrm{keV}$, this is still large, but the excited states will have a non-negligible occupation number that can be calculated from a Boltzmann distribution. It is easy to acknowledge that the amount of axions produced is proportional to the occupation number, the isotope abundance and the inverse lifetime of the excited state. By combining a list of possible elements and their nuclear transitions [373] with solar abundances [374], it becomes clear that for IAXO the lowest transition of ${}^{57}\mathrm{Fe}$ would produce the strongest signal (see appendix F for more details). In the following, this will be the only isotope under consideration, but the calculation would be completely analogous for all other possible transitions.

Axion-nucleon interactions were defined in the derivative basis in equation (7.1). For calculating nuclear transition rates, they are commonly rewritten as

$$\mathcal{L}_{aN} = -ia\bar{N}\gamma_5\left(g_{aN}^0 + g_{aN}^3\tau^3\right)N\ . \tag{7.40}$$

$N = (p,n)^T$ denotes the nucleon doublet and g_{aN}^0 and g_{aN}^3 are the iso-scalar and iso-vector couplings, respectively. τ^3 is the third Pauli matrix. Note that the transition from derivative to Yukawa-type interactions is justified as long as the difference in anomalous gauge boson couplings does not enter the calculation.

The axion-to-photon branching ratio for the decay rates of the first excited state of ${}^{57}\mathrm{Fe}$ can then be expressed as [375, 376]

$$\frac{\Gamma_a}{\Gamma_\gamma} = \left(\frac{k_a}{k_\gamma}\right)^3 \frac{1}{2\pi\alpha_{\mathrm{EM}}}\frac{1}{1+\delta^2}\left[\frac{\beta\, g_{aN}^0 + g_{aN}^3}{\left(\mu_0 - \frac{1}{2}\right)\beta + \mu_3 - \eta}\right]^2 , \tag{7.41}$$

where k_a, k_γ are the axion and photon momenta, μ_0 and μ_3 are the isoscalar and isovector nuclear magnetic moments (expressed in nuclear magnetons), δ is the E2/M1 mixing ratio for the ${}^{57}\mathrm{Fe}$ nuclear transition and β and η are constants dependent on the nuclear structure [376]. The values of the isotope dependent constants were reevaluated in ref. [377]. Using these and taking the relativistic limit $(k_a/k_\gamma) \to 1$, the branching ratio becomes

$$\frac{\Gamma_a}{\Gamma_\gamma} = 2.32\left(-1.31\, g_{aN}^0 + g_{aN}^3\right)^2 \tag{7.42}$$

$$= 2.32\left(0.16\, g_{ap} + 1.16\, g_{an}\right)^2\ . \tag{7.43}$$

The second line reinstates the coupling constants as defined in equation (7.1) by using the identities

$$g_{an} = g^0_{aN} - g^3_{aN} \,, \tag{7.44}$$

$$g_{ap} = g^0_{aN} + g^3_{aN} \,. \tag{7.45}$$

It is convenient to define the combination appearing in the branching ratio as an effective axion-nucleon coupling,

$$g^{\text{eff}}_{aN} = 0.16\, g_{ap} + 1.16\, g_{an} \,, \tag{7.46}$$

which is specific to this transition. Notably, the updated branching ratio cited here is 27 % larger than the one found in ref. [375], which was used in the analyses of refs. [203, 212, 378].

From the branching ratio, the axion emission rate per unit of solar mass can be computed as

$$\mathcal{N}_a = \mathcal{N}\, \omega_1(T)\, \frac{1}{\tau_0} \frac{1}{1+\alpha} \frac{\Gamma_a}{\Gamma_\gamma}, \tag{7.47}$$

where $\mathcal{N}$ is the ^{57}Fe number density per solar mass, ω_1 the occupation number of the first excited state, τ_0 the lifetime of the excited state and α the internal conversion coefficient. ω_1 is computed from a Boltzmann distribution as in ref. [203] or in appendix F. The natural linewidth of the emission peak $\sim 1/\tau_0$ is negligible compared to the Doppler broadening σ given by

$$\sigma(T) = E^* \sqrt{\frac{T}{m_{\text{Fe57}}}} \sim 2\,\text{eV} \,. \tag{7.48}$$

This corresponds to a full width at half maximum (FWHM) of $2.35\,\sigma \sim 5\,\text{eV}$. Using this as the spectral shape, one can find the axion emission from the ^{57}Fe transition per axion phase space Γ^{Fe57}_a as

$$\Gamma^{\text{Fe57}}_a(\omega) = \frac{(2\pi)^3}{4\pi\omega^2}\, \mathcal{N}_a\, \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(\omega - E^*)^2}{2\sigma^2}\right) \rho \,. \tag{7.49}$$

The first factor is required for the normalization of the ω integral, which will become clear in the context of section 7.1.7. The quantities $\mathcal{N}_a$, ρ and σ are all functions of the solar radius. One would not expect the Doppler broadened peak to be resolved in most detection scenarios in which case only the total flux is of interest. The integral over ω can then be performed analytically and only the one over the solar radius remains to be evaluated numerically.

Note that this axion source is particularly sensitive to the solar temperature. The reason for this is that the thermal occupation number $\omega_1 \propto e^{-E^*/T}$ of the excited nuclear state is highly suppressed. For instance, the latest high- and low-metallicity solar models (B16-GS98 and B16-AGSS09 [379]) only differ by about 1 % in their respective core temperature, but this alone results in a difference of 12 % in the total flux.

Because the 14.4 keV line from ^{57}Fe is well separated from the rest fo the spectrum, it is left aside in the following discussions of uncertainties and detection opportunities of helioscopes, which are optimised for 1 keV to 10 keV. Instead, section 7.5 is entirely focused on the kinds of experimental setups which could be sensitive to axions from nuclear interactions.

7.1.6 Electron densities and ionisation states

All of the production rates derived in the previous sections depend on a small set of solar plasma parameters. Besides the monochromatic opacity and the solar magnetic field, which require extra attention, these are the temperature T, the number density of free electrons n_e and the number density of ions weighted by their charge squared $\bar{n}$. Publicly available standard solar models only contain the values of T at every radius, as well as other solar properties like density, pressure and element abundances. The values of n_e and $\bar{n}$ are usually not tabulated and hence need to be computed.

Using the information provided by the standard solar model format, there are two ways to calculate n_e. The first option is to neglect all elements heavier than helium and to assume full ionisation. Under these assumptions, the electron density can be computed from the matter density ρ and the mass fraction of hydrogen X as[49] [231]

$$n_e \approx \frac{X+1}{2} \frac{\rho}{m_u} \, , \tag{7.50}$$

where m_u denotes the atomic mass unit. The reason why this approximation works very well is that it is exact for fully ionised hydrogen and helium, which together account for more than 98 % of the total solar mass. Furthermore, the remaining heavier elements also contribute at most one electron per two nucleons at full ionisation, which is how they enter in the expression above.

The other option is to explicitly sum over all elements in the Sun and directly assume full ionisation,

$$n_e \approx \sum_z z \, n_z = \sum_z z \, X_z \, \frac{\rho}{A_z m_u} \; . \tag{7.51}$$

In this equation, X_z denotes the mass fraction of each element and A_z is its standard atomic weight. The resulting electron density only differs from the one computed with equation (7.50) on the sub-percent level.

By construction, both of the approximations above must return a larger value than the true one. For this reason, all following computations use the smaller value, which is the one summing over all elements and assuming full ionisation in equation (7.51).

The next parameter $\bar{n}$ appears both in the Primakoff rate (7.9) as well as in the Debye screening scale (7.13). To calculate it, Raffelt used the same approximation

---

[49] This thesis adopts the commonly used notation of writing the mass fraction of all hydrogen isotopes as X and the mass fraction of all helium isotopes as Y, which in turn defines the metallicity $Z \equiv 1 - X - Y$.

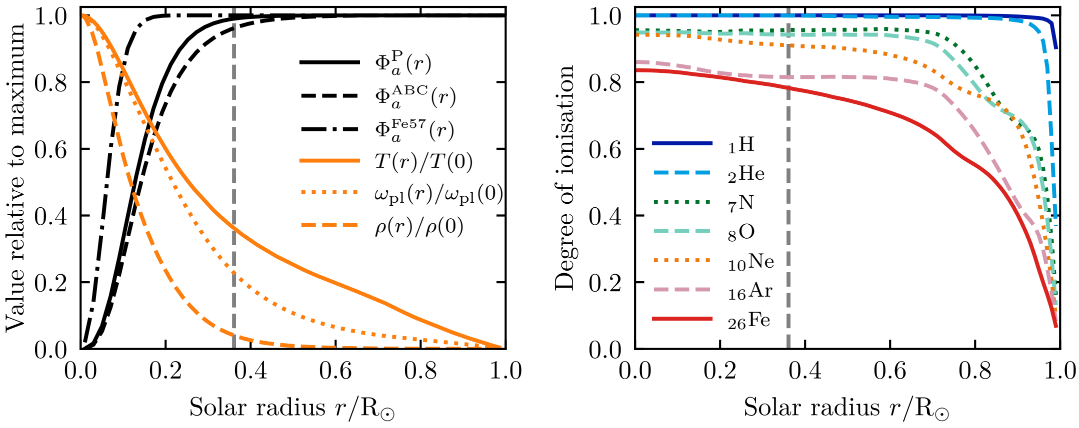

Figure 7.5. Radial dependence of various solar quantities. The solar model B16-AGSS09 was used for this plot [379]. The vertical, dashed grey lines in both panels correspond to the radius below which 99 % of the total Primakoff and ABC fluxes are generated. *Left panel:* The Primakoff (Φ_a^{P}) and ABC fluxes (Φ_a^{ABC}) integrated to the given solar radius as well as the solar temperature (T), density (ρ) and plasma frequency (ω_{pl}) with respect to their maximum values. *Right panel:* Degree of ionisation, Q_z/z, or various elements in the Sun. The values were extracted from the OP code [365, 380].

as the one for n_e in equation (7.50) [231]. This simply results in

$$\bar{n} = \sum_z Q_z^2 \, n_z = \sum_z Q_z^2 \, X_z \, \frac{\rho}{A_z \, m_u} \tag{7.52}$$

$$\approx \frac{\rho}{m_u} \, , \tag{7.53}$$

where only the first two terms for hydrogen and helium (for which $Q_z^2/A_z = 1$) were included in the second line and the metallicity was approximated with zero, i.e. $X + Y = 1$. This approximation is however significantly less accurate for $\bar{n}$ than it is for n_e. The reasons is that due to the square of charges, heavy elements can contribute significantly to the total value of $\bar{n}$ even though they only make up a small fraction of the total density. Despite these heavy ions having large charges, full ionisation is also not a good approximation. This can be intuitively seen by comparing the full ionisation energies of heavy elements with the core temperature. The former far exceeds the latter and it should hence not be expected that metals are fully ionised in the solar plasma.

Because both of the previously successful approximations fail to give $\bar{n}$ with sufficient accuracy, the expectation value of Q_z^2 has to be computed for each element at every solar radius. This can be done using data from the Opacity Project (OP) [380]. Besides calculating opacities for arbitrary chemical compositions of solar-like plasmas, the OP also provides tables of ionisation states for heavy ions in such plasmas. These tables can be interpolated to calculate Q_z^2 as a function of T and n_e. The results for some relevant elements are plotted in the right panel of figure 7.5. Note that due to the dependence of the ionisation on n_e, the same method could not be used for calculating n_e itself. It is however possible to check the consistency of the calculation by comparing n_e from OP data (using equation (7.51) as input) with the full ionisation result. The relative difference does not exceed 0.2 % in the axion production region. This is small enough for all current practical purposes and it proves the consistency of the values adopted for n_e and $\bar{n}$.

Figure 7.5 illustrates the radial dependence of some important solar values as well as the accumulation of the total axion flux inside the Sun. It is important to note that most axions are generated inside the solar core ($r/\mathrm{R}_\odot < 0.25$) and almost the entire flux is emitted until about $r/\mathrm{R}_\odot = 0.5$. The three flux components shown in figure 7.5 are sorted by their energy. The 14.4 keV axions from ^{57}Fe transitions are only generated extremely close to the centre. Primakoff and ABC processes also happen further outwards but it can be seen that the harder Primakoff spectrum is produced more closely to the centre than the ABC flux. This observation is further support for the assumption of full ionisation in the calculation of n_e because this approximation is better at high temperatures in the deep solar interior.

7.1.7 Integration of production rates over the solar model

Helioscopes can track and observe the whole or parts of the Sun throughout the day. An accurate signal prediction requires a calculation of the spectral axion flux generated inside the field of view. Because the distance of the Sun to Earth d_E is much larger than the solar radius, $\mathrm{R}_\odot/d_\mathrm{E} \approx 0.005$, it is a good approximation to assume that all axions are produced at the same distance from the experiment and that parallax is negligible [202]. Using this, the total spectral flux $\mathrm{d}\Phi_a/\,\mathrm{d}\omega$ can be written as an integral of the production rate Γ over the solar volume V [105],

$$\frac{\mathrm{d}\Phi_a}{\mathrm{d}\omega} = \frac{1}{4\pi d_\mathrm{E}^2} \int \mathrm{d}V \, \frac{4\pi\omega^2}{(2\pi)^3} \, \Gamma(\omega, r) \; , \tag{7.54}$$

where the first factor in the integrand accounts for the directional part of the phase space integral over axion momenta (still assuming relativistic axions). d_E is the (averaged) distance between Sun and Earth during the observation[50] and Γ can be any (or the sum of all) of the production rates discussed earlier in this section.

In a spherically symmetric solar model, the integral becomes one-dimensional[51]

$$\begin{aligned} \frac{\mathrm{d}\Phi_a}{\mathrm{d}\omega} &= \frac{\mathrm{R}_\odot^3}{4\pi d_\mathrm{E}^2} \frac{4\pi\omega^2}{(2\pi)^3} \int_0^{2\pi} \mathrm{d}\phi \int_0^{\pi} \mathrm{d}\theta \, \sin\theta \int_0^1 \mathrm{d}r \, r^2 \, \Gamma(\omega, r) && (7.55) \\ &= \frac{\mathrm{R}_\odot^3 \omega^2}{2\pi^2 d_E^2} \int_0^1 \mathrm{d}r \, r^2 \, \Gamma(\omega, r) \; , && (7.56) \end{aligned}$$

where the radius r is expressed in units of the solar radius $\mathrm{R}_\odot$.

Depending on the optics used in a helioscope, it could be possible to resolve different regions of the solar disc. As an example, one can consider a circular central part of the solar disc with radius ρ_1. In this case, it is convenient to perform the

---

[50] Throughout this work, all fluxes are computed with $d_\mathrm{E} = 1\,\mathrm{au} \approx 1.496 \times 10^{11}\,\mathrm{m}$ [381].

[51] The macroscopic magnetic field in equation (7.28) is not spherically symmetric. Hence, the projected magnetic fields $B_\perp$ and $B_\parallel$ have to be averaged over plasmon directions and the spherical shell inside the Sun, as it was done in refs. [347, 356, 357].

integral in cylindrical coordinates, where $r^2 = \rho^2 + z^2$ and $|z| < z_{\mathrm{max}} = \sqrt{1-\rho^2}$ defines the solar sphere. The volume integral can then be written as

$$\frac{\mathrm{d}\Phi_a}{\mathrm{d}\omega} = \frac{\mathrm{R}_\odot^3 \omega^2}{(2\pi)^3 d_\mathrm{E}^2} \int_0^{2\pi} \mathrm{d}\phi \int_0^{\rho_1} \mathrm{d}\rho\, \rho \int_{-z_{\mathrm{max}}(\rho)}^{z_{\mathrm{max}}(\rho)} \mathrm{d}z(\rho)\; \Gamma(\omega, r) \tag{7.57}$$

$$= \frac{\mathrm{R}_\odot^3 \omega^2}{(2\pi)^2 d_\mathrm{E}^2} \int_0^{\rho_1} \mathrm{d}\rho\, \rho \left(2 \int_\rho^1 \mathrm{d}r \left| \frac{\mathrm{d}z}{\mathrm{d}r} \right| \Gamma(\omega, r) \right) \tag{7.58}$$

$$= \frac{\mathrm{R}_\odot^3 \omega^2}{2\pi^2 d_\mathrm{E}^2} \int_0^{\rho_1} \mathrm{d}\rho\, \rho \int_\rho^1 \mathrm{d}r\, \frac{r}{\sqrt{r^2 - \rho^2}}\, \Gamma(\omega, r)\;. \tag{7.59}$$

In principle, a good resolution of this geometry with respect to ρ_1 can give access to the radial dependence of the emission rate $\Gamma(\omega, r)$.

7.2 Solar models and opacity codes

Solar models are the key input for calculating the solar axion flux. In this section, all of the publicly available models with sufficient information to compute the production rates are collected and qualitatively compared. This is followed by a more in-depth quantitative analysis of associated uncertainties in the next section. Table 7.1 provides an overview of the eleven models under consideration. The first noteworthy difference is that some provide less information than others on the abundances of heavier solar metals, e.g. Fe, which is of particular relevance for the calculation of the ABC flux. Furthermore, all of the solar models available today fail to consistently explain both seismological and photospheric measurements. This is known as the solar metallicity or abundance problem [374, 382–386] in the literature and it has so far not been fully resolved [387]. Table 7.1 clearly illustrates this by showing whether a given model is in agreement with either seismological observations or photospheric data. Note also that this correlates with the overall metallicity of the model. Models with larger values of Z tend to agree with helioseismology, while low-Z models fit better to photospheric abundances.

The differences between solar models, has significant impact on the resulting axion spectrum. This is shown in figure 7.6, where all fluxes except the 14.4 keV line[52] are shown for all solar models listed in table 7.1. Since the BP and BS models do not contain information on the abundance of heavier metals, the flux from axion-electron interactions (Φ_a^{ABC}) is missing contributions from these heavy constituents. This results in the absence of peaks caused by atomic transitions and a reduction of the smooth part of the spectrum by about 20–40 %. In the case of the Primakoff spectrum, the deviations between BP/BS models and all other models enter through differences in the Debye screening scale. The spectral flux can be up to about 12.5 % *larger* than the Primakoff flux from e.g. the B16-AGSS09 model. While the differences are not as striking as they were for the ABC flux, the neglect of heavier metals still causes a relatively large systematic difference. Since all of these deviations are entirely due to a lack of information, it is clear that these models are less suited for solar axion flux evaluations.

---

[52] This monochromatic line will be discussed in detail in section 7.5.

Solar model	**Comments**	**S**	**Ph**	Z	**References**
BP98	no heavier metals tracked	✓	✗	0.0201	[389][a]
BP00	no heavier metals tracked	✓	✗	0.0188	[390][a]
BP04	no heavier metals tracked	✓	✗	0.0188	[391][a]
BS05-OP	no heavier metals tracked	✓	✗	0.0190	[392][a]
BS05-AGSOP	no heavier metals tracked	✗	(✓)	0.0141	
AGS05		✗	✓	0.0140	[385]
GS98		✓	✗	0.0188	[385, 388][b]
AGSS09	same as AGSS09met	✗	✓	0.0150	
AGSS09ph		✗	✓	0.0152	
B16-GS98	updated version of GS98	✓	✗	0.0188	[379][c]
B16-AGSS09	updated version of AGSS09(met)	✗	✓	0.0150	

[a] Model files available at www.sns.ias.edu/~jnb/SNdata/solarmodels.html
[b] Model files available at wwwmpa.mpa-garching.mpg.de/~aldos/solar_main.html
[c] Model files available at www.ice.csic.es/personal/aldos/Solar_Data.html

Table 7.1. List of the solar models considered in this work in order of their respective publication date. Green ticks and red crosses denote, respectively, whether or not the models (roughly) agree with helioseismological ("S"; high Z) or photospheric measurements ("Ph"; low Z). In BP and BS models only a few metal abundances are tabulated. These can therefore only agree with the photospheric measurements of these particular metals, which is marked by a tick in parentheses. Z denotes the total mass fractions of metals, which is larger than the surface metallicity (see refs. [385, 388]) due to diffusion and gravitational settling [374].

In order to get a better idea of what the true systematic differences between competing solar models is, figure 7.7 shows the deviations between models which provide the most complete information on individual isotope abundances. Two of them (GS98 and B16-GS98) fall in the high-Z category, while the other three have a significantly lower values of Z. The deviations in the Primakoff flux are typically a few, but no more than about 10 % and are most pronounced in the high energy tail of the spectrum. The ABC flux shows a similar trend with deviations up to about 15 %. In addition, the transition peaks are subject to particularly large deviations. In conclusion, the choice of solar model has crucial influence on the solar axion flux and systematic uncertainties related to the competing solar models are more pronounced for the ABC than for the Primakoff flux.

In ref. [105] a powerful relation between the ABC axion flux and the monochromatic opacity was derived, which is reviewed in section 7.1.4. The required values of the solar opacity can be computed using one of a number of available opacity codes. Table 7.2 lists the ones considered in this thesis. The only code which is easily adoptable to arbitrary chemical compositions is the one by the opacity project (OP). LEDCOP and ATOMIC also work for arbitrary mixtures but they can only

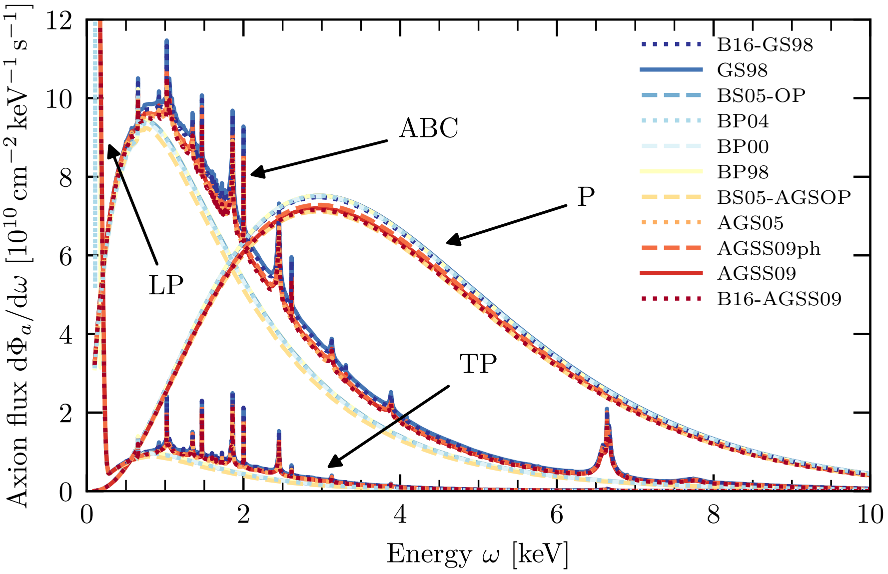

Figure 7.6. Solar axion fluxes for different solar models between 0.1 keV and 10 keV, using opacities from the OP code. Coupling constants were chosen as $g_{ae} = 10^{-12}$ and $g_{a\gamma\gamma} = 10^{-10}\,\mathrm{GeV}^{-1}$.

be generated via an online interface, which is unfeasible for the very large number of different models that is considered in the next section. Finally, the OPAS collaboration only provided monochromatic opacities for the AGSS09 model, but these are not publicly available.

Following the same logic as before, figure 7.8 summarises how different opacity codes affect the ABC flux, using AGSS09 as a reference solar model.[53] At first sight in the left panel of figure 7.8, the spectral flux agrees rather well – especially above ∼2 keV. Nevertheless, the relative differences are found as typically 10–20 % across the whole energy range and there are additional huge deviations around some of the peaks in the spectrum. The latter are particularly important for studying the metal content of the Sun with suitable energy resolution [346] but become less of an issue for integrated flux measurements. As pointed out in ref. [105], the OPAS code is not fully applicable to the axion calculation below ∼2 keV. This is why the resulting spectrum is also only plotted up to this energy in figure 7.8. Because of its superior flexibility, the OP code is used in all of the following calculations and the systematic uncertainties on the continuous part of the spectrum are quantified in section 7.3. It should however be kept in mind that a different choice of opacity code can significantly affect the outcome of a solar axion study, especially when peaks from atomic transitions are investigated.

---

[53] Because the OPAS only provided data for the AGSS09 composition, this is the only possible choice for a consistent comparison.

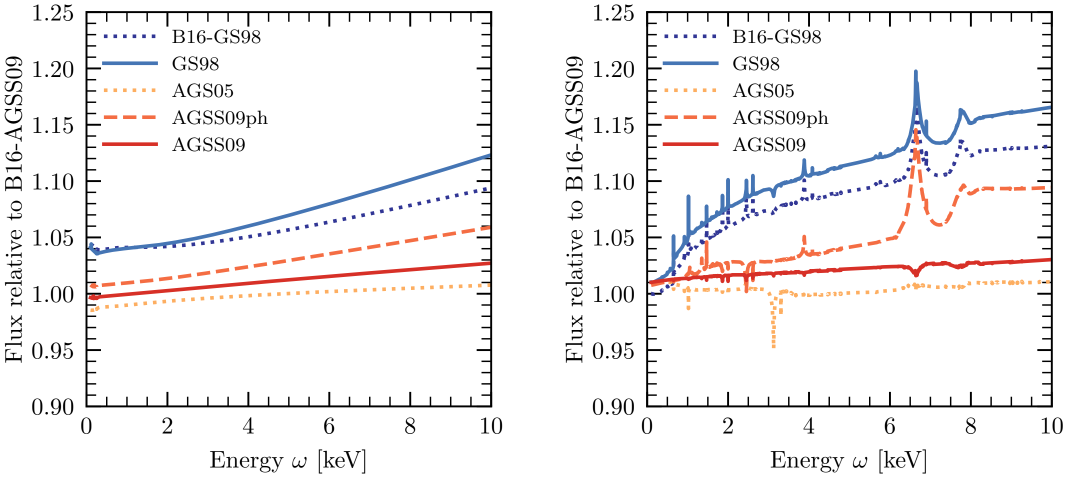

Figure 7.7. Relative deviations in the axion flux from different solar models (compared to the B16-AGSS09 model) for axion-photon interactions (left panel) and axion-electron interactions (right panel).

Opacity codes	Comments	Arb. comp.	References
OP		✓	[365, 380]
LEDCOP		(✓)	[366][a]
ATOMIC		(✓)	[393][a]
OPAS	Only for AGSS09 comp.	✗	[367, 394][b]

[a] Available at https://aphysics2.lanl.gov/apps/ .

[b] The monochromatic opacities for the AGSS09 composition are not publicly available but were provided by C. Blancard and the OPAS collaboration.

Table 7.2. List of opacity codes considered in this work in order of their respective publication date. Green ticks and red crosses denote whether or not it is possible to calculate opacities for arbitrary solar composition. With the LEDCOP and ATOMIC codes, this is only possible via an online interface.

7.3 Uncertainties of the solar axion flux

While it is interesting to compare how the solar axion flux differs for various solar models and opacity codes, it is not sufficient as a complete representation of the total uncertainty. Instead, it is necessary to investigate the intrinsic, statistical uncertainties associated with a given combination of solar model and opacity code. Because only the OP code allows for full flexibility regarding different solar compositions, it is the natural choice as a benchmark for the opacity calculation.

The first two parts of this section focus on the uncertainties arising from the modelling of the solar plasma properties. Systematic uncertainties of the calculation are discussed in section 7.3.4. At all times the results come from the Primakoff rate as in equation (7.7) including the electron degeneracy factor (7.15) and from the full ABC flux (without Pauli blocking) (7.38). For simplicity, the flux from plasmon

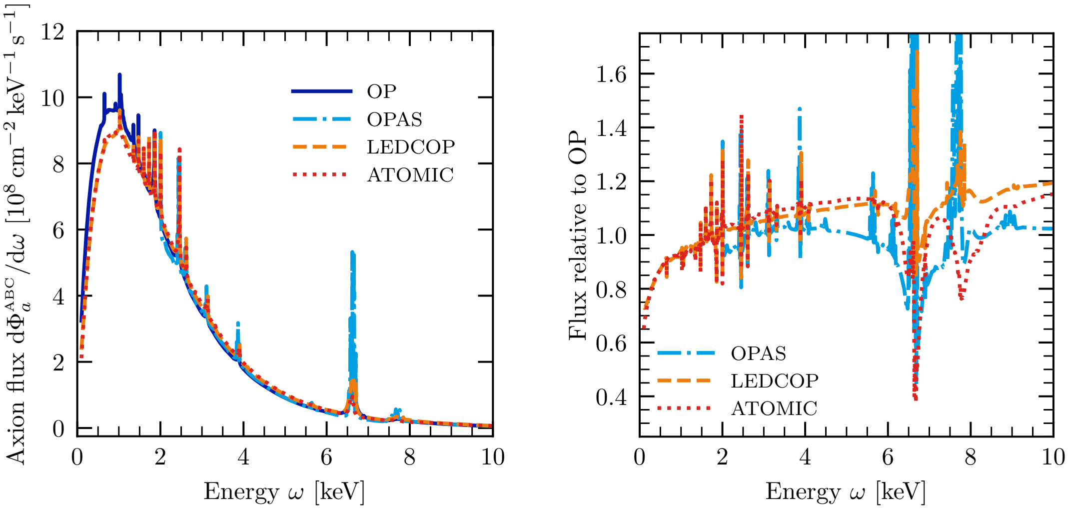

Figure 7.8. ABC flux (left panel) and relative deviations compared to the OP code (right panel) for the AGSS09 model and different opacity codes. The low-energy region below 2 keV is excluded for OPAS because this is where this code may not be fully accurate [105]. In the right panel, the huge deviations (up to 440 %) for the Fe peak around 6.5 keV are also not fully shown.

conversions and nuclear transitions are not included. This is because the plasma conversion is much more strongly affected by the uncertainties in the modelling of the solar magnetic fields and it is anyway expected to be subdominant in the 1 keV to 10 keV range, see figure 7.4. The ^{57}Fe line also lies outside of this range and its phenomenological impact is separately discussed in section 7.5.

7.3.1 Solar model uncertainties

To explore how uncertainties of solar models translate into uncertainties of the solar axion flux, one can exploit results produced by the authors of ref. [395]. They generated representative samples of about 10,000 solar models from a Monte Carlo (MC) simulation. This was done using the Garching stellar evolution code [396]. It can be used to evolve an initially homogeneous star over the full lifetime of the Sun. By varying three free parameters (initial helium abundance, initial metallicity and mixing length parameter), one can ensure that the model converges to the desired present solar properties (luminosity, radius and chemical composition at the surface) within a relative accuracy of 10^{-4}, ending up with a realistic model for the current state of the Sun. Apart from the three free parameters, which are varied in order to fulfill the criteria of convergence, the procedure depends on a total of 21 input parameters, which can have a strong effect on the final solar model. These include rates of nuclear fusion reactions, solar properties (e.g. luminosity, age and diffusion coefficient), chemical composition as well as opacities and equations of state. All of these inputs come with their own uncertainty, which can be expressed in the form of a set of independent probability distributions, one for each input parameter. For the 10,000 solar models in ref. [395] the input parameters are drawn from these probability distributions, before applying the procedure summarised above. This method ensures the generation of a statistically representative sample of solar models.

As mentioned earlier, the solar abundance problem [374, 382–386] currently still persists, with models from the high-Z as well as the low-Z category being favoured by helioseismology and photospheric observations, respectively. To cover both extremes, sets of 10,000 models are considered for one representative of each of the two categories, namely AGSS09 (low Z) and GS98 (high Z).

The original idea of ref. [395] was to use this sample to quantify uncertainties of the solar neutrino fluxes, helioseismological quantities and other characteristic values such as the core temperature. Clearly, the approach is very flexible and can readily be applied to solar axion fluxes. For this, updated sets of each of the AGSS09 and GS98 solar model were kindly provided by the authors of ref. [395]. By calculating all of the corresponding solar axion spectra as described in section 7.1, the statistical uncertainty of the axion flux can be directly inferred from the statistical uncertainties of the input parameters used to build a solar model. The MC samples for solar models automatically take into account all possible dependencies between the solar parameters that enter the axion flux calculation.

7.3.2 Opacity code uncertainties

In addition to the solar model, the ABC flux crucially depends on the adopted opacities. This was already demonstrated by the direct comparison in figure 7.8. Some of these large deviations can however be identified as unphysical [105] and it would be desirable to have a more robust quantification of the actual uncertainties.

The benchmark opacity code throughout this work is the one by the OP because it allows to compute opacities for arbitrary chemical compositions and therefore equation (7.31) can be used for all ff processes involving H and He. Furthermore, it is possible, using the OP code, to calculate opacities for each mixture of the $2 \times 10,000$ solar models individually. Most importantly, the OP does not make assumptions or include effects that are fundamentally not applicable to axion production [105]. However, it does not supply its own uncertainty estimates of monochromatic opacities. In order to quantify these, one can exploit recent efforts in solar modelling, in which uncertainties of the opacity input was carefully analysed [379]. It was argued that opacity uncertainties are much larger at the bottom of the convective zone (CZ) compared to the solar centre. This led ref. [379] to a temperature-dependent relative opacity uncertainty $\delta\kappa(T)$, which can easily be adopted for the MC simulations of this work. As suggested, a logarithmic interpolation between the uncertainty level at the core (2 %) and at the bottom of the convection zone (7 %) is used. This means the opacity from the OP κ_{OP} is modified as

$$\kappa(\omega;\, T,\, n_e) = \kappa_{\mathrm{OP}}(\omega;\, T,\, n_e)\ \left(1 + \delta\kappa(T)\right) \, , \tag{7.60}$$

where

$$\delta\kappa(T) = a + b\,\frac{\log_{10}\left(T_0/T\right)}{\log_{10}\left(T_0/T_{\mathrm{CZ}}\right)} \, , \tag{7.61}$$

with $\log_{10}(T_0/T_{\rm CZ}) \simeq 0.9$ and where $T_0 \equiv T(r = 0)$ is directly taken from the solar model. For the MC simulation, a and b are drawn from normal distributions with $\sigma_a = 0.020$ and $\sigma_b = 0.067$ [379], which reproduces the uncertainty levels in the core and at the bottom of the convective zone.

The procedure above can directly be incorporated into the MC simulation. However, this applications differs from the one in solar modelling because it is applied to monochromatic opacities. In solar modelling, on the other hand, one is mostly interested in radiative energy transport, which is typically quantified by frequency averaged opacities like the Rosseland mean (see equation (7.25)). One may therefore wonder why the uncertainty estimation of the averaged opacity should be equally applicable to the monochromatic one. In support of this, it can be argued that the spectral shape of the opacity is also varied in the approach used here. This is not done via equation (7.60) but through the variation of the chemical composition within the set of 10,000 solar models. The full MC procedure does therefore not merely rescale $\kappa(\omega)$ by a constant factor. Furthermore, it can be expected that the method used here results in conservatively large uncertainties rather than an underestimation. The reason for this is that a large fraction of the opacity uncertainty cited in ref. [379] actually stems from the uncertainty of the chemical composition. By varying both composition and overall opacity, one is therefore more likely to overestimate the total uncertainty.

These arguments only apply to the continuous part of the spectrum, where it is very reasonable to assume that the uncertainty of the monochromatic opacity contribution from one element is not much larger than the uncertainty of the Rosseland mean. As shown in figure 7.8, the various opacity codes produce widely different values at the peaks from atomic transitions. This has little to no influence on the resulting frequency averaged opacities and it is therefore not included in the uncertainty estimation above. This means that one can still only estimate the uncertainty in the solar axion flux around the peaks by comparing different opacity codes. Results from the MC simulation in the following section should therefore not be understood as a complete estimation of the uncertainty at these transition energies. This also has consequences for solar axion searches: even if the peaks can be experimentally resolved, the uncertainty in the opacity calculation causes problems for abundance measurements of solar metals [346].

7.3.3 Results of the Monte Carlo simulation

Having established the MC procedure for both solar models and opacity evaluations, the spectral axion fluxes can be calculated. This is done for a precise total of 9,979 realisation of AGSS09 and 9,971 of GS98. These are all of the models provided by the authors of ref. [395]. The results are shown in figure 7.9, where the relative differences of the two models are shown together with the 1 σ deviation. Just like the comparison in figures 7.7 and 7.8, the fluctuations are larger at higher energies. The relative uncertainty is however significantly smaller and only reaches a few percent. As one would expect from the greater complexity and higher number of processes involved, the typical uncertainties for the ABC spectrum are generally larger than for the Primakoff flux.

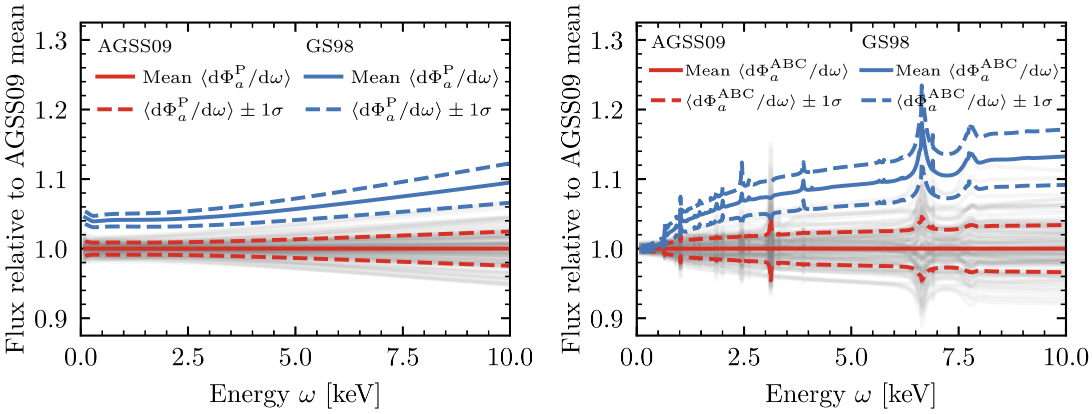

Figure 7.9. Comparison of solar axion spectra calculated from the Monte Carlo simulation. All values are plotted relative to the mean of the AGSS09 model. Mean values of AGSS09 (red) and GS98 (blue) are shown as solid lines and $\pm 1\sigma$ bands as dashed ones. The transparent grey lines are 100 randomly chosen Monte Carlo samples of the AGSS09 model. Primakoff and ABC fluxes are shown in the left and right panel, respectively.

The total axion flux will in general be comprised of both Primakoff and ABC spectra at an arbitrary ratio. To visualise how this affects the total uncertainty, it is helpful to define the flux ratio γ as

$$\gamma \equiv \frac{\Phi_a^\mathrm{P}}{\Phi_a^\mathrm{ABC}} = 1.32 \left(\frac{g_{a\gamma\gamma}}{10^{-10}\,\mathrm{GeV}^{-1}} \right)^2 \left(\frac{10^{-12}}{g_{ae}} \right)^2 . \tag{7.62}$$

The numerical prefactor is obtained by integrating the spectrum between 0.1 keV and 10 keV for the AGSS09 model and OP opacities. $\gamma \gg 1$ or $\gamma \ll 1$ corresponds to the domination of Primakoff or ABC flux, respectively.

Figure 7.10 shows the relative uncertainties of the spectral flux as a function of γ. Typical values are of the order of a few percent. Significantly larger uncertainties only occur at the peaks in the ABC spectrum, where the systematic uncertainties are larger anyway. In the limit of large or small γ, the values shown in figure 7.9 are recovered.

7.3.4 Discussion of additional uncertainties

The MC simulation above is able to account for the impact of statistical uncertainties of the input data and parameters. There are, however, further systematic uncertainties and corrections, which originate from the approximations employed in the calculations of section 7.1. All of these corrections considered in this work are listed in table 7.3 alongside the MC results for comparison.

Primakoff flux

Besides the Primakoff effect, there are other production processes based on the axion-photon coupling which were neglected above. These comprise Compton scattering, bremsstrahlung or atomic interactions which emit a photon and an axion instead of

Correction/Uncertainty	**Order**	**Averaged**	**Maximal**
Primakoff flux			
Solar model uncertainty (systematic)		$\sim 5.1\,\%$	$\sim 11\,\%$
Solar model uncertainty (statistical)		$\sim 1\,\%$	$\sim 2.5\,\%$
Atomic transition (ff, bf, & bb) emitting $\gamma + a$		$< 0.2\,\%$	$\sim 4\,\%$
Higher-order QED effects	α_{EM}	$< 0.7\,\%$	$< 0.7\,\%$
Electro-Primakoff effect		$< 4 \times 10^{-5}$	$< 0.4\,\%$
Non-vanishing axion mass ($m_a \sim$ eV)	m_a/ω	$< 0.1\,\%$	$< 0.1\,\%$
Inelastic Primakoff		$\lesssim 0.1\,\%$	$\lesssim 0.1\,\%$
Form factor for non-static charges		$\lesssim 0.02\,\%$	$\lesssim 0.02\,\%$
Full relativistic dispersion relation		$\lesssim 5 \times 10^{-6}$	$\lesssim 10^{-6}$
Non-resonant transverse plasmon conversion†		$0.01\,\%$–$5\,\%$	$0.02\,\%$–$50\,\%$
Resonant longitudinal plasmon conversion†		$< 2 \times 10^{-6}$	$< 0.7 \times 10^{-4}$
Resonant longitudinal plasmon conversion			
Solar magnetic field strength		factor 225	factor 225
Non-resonant transverse plasmon conversion			
Solar magnetic field strength		factor 225	factor 225
Total scattering rate Γ_{T} (Rosseland mean vs. monochromatic)		factor 2.2	factor 15.2
ABC flux			
Solar model uncertainty (systematic)		$\sim 5.4\,\%$	$\sim 19\,\%$
Solar model uncertainty (statistical)		$\sim 1.5\,\%$	$\sim 5\,\%$
Opacity uncertainty (systematic)		$1\,\%$–$3\,\%$	$< 440\,\%$
OP opacity uncertainty (statistical)		$\lesssim 1\,\%$	$\sim 17\,\%$
Born approximation in eqs. (7.31) & (7.32)		$\lesssim 10\,\%$	$\lesssim 10\,\%$
Electron degeneracy		$\lesssim 7\,\%$	$\lesssim 8\,\%$
Higher-order in multipole expansion in equation (7.37)	$\omega/(m_e Z \alpha_{\text{EM}})$	$\lesssim 2\,\%$	$\lesssim 7\,\%$
Effects of relativistic electron		$\lesssim 3\,\%$	$\lesssim 6\,\%$
Higher-order QED effects	α_{EM}	$< 0.7\,\%$	$< 0.7\,\%$
Spatial & spin wave function non-separable	$E_n\,(Z\alpha_{\text{EM}})^2/n$	$\lesssim 0.5\,\%$	$\lesssim 0.5\,\%$*
Non-vanishing axion mass ($m_a \sim$ eV)	m_a/ω	$< 0.1\,\%$	$< 0.1\,\%$

† Plasmon fluxes are distinct and experimentally distinguishable from the Primakoff flux (see section 7.1.3) and are therefore not strictly speaking corrections of the Primakoff flux. They are only included here to quantify their potential relevance.

* The assumption of separability of spatial and spin wave functions only works if the atomic fine structure is not resolved. In practice this limits the resolution to about $\sim 25\,$eV [397]. The uncertainty given here corresponds to a spectrum smoothed out at this level of resolution. With better resolution, the local uncertainty can be much larger.

Table 7.3. Corrections and uncertainties relative to the respective solar axion flux. Where possible, a parametric estimate is provided next to the averaged and maximal effect between 1 keV and 10 keV.

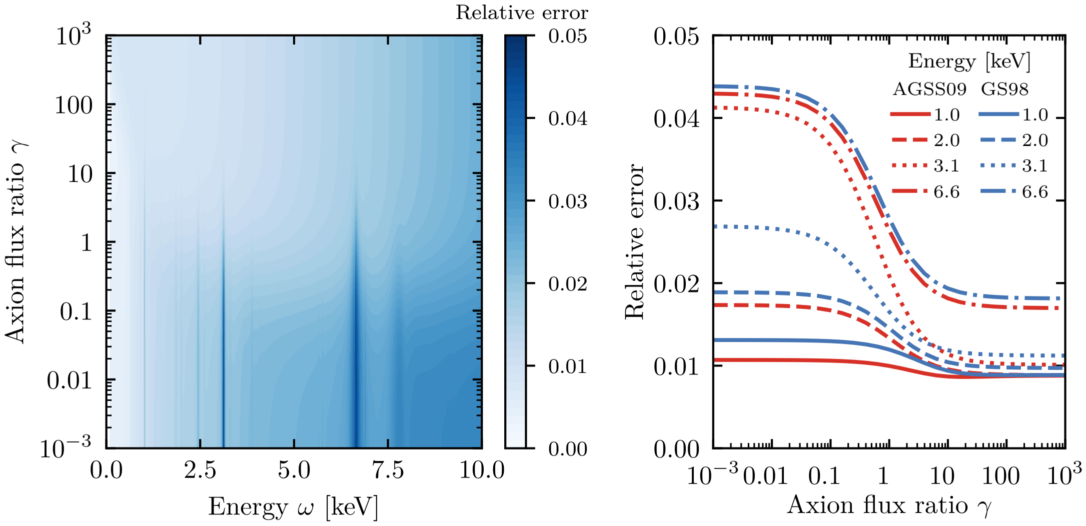

Figure 7.10. Relative errors of the solar axion flux as a function of energy ω and flux ratio γ. The density plot on the left is for the AGSS09 solar model. In the right panel, both AGSS09 (red lines) and GS98 (blue lines) are included. Note that the systematic uncertainties around the peaks can be notably larger than a few percent (see figure 7.8).

only a photon. The corresponding rates can in principle be calculated and included but they come out as highly suppressed. This can be seen from an approximate expression analogous to equation (7.37),

$$\frac{\Gamma^i_{\gamma+a}}{\Gamma^i_\gamma} \sim g^2_{a\gamma\gamma}\,\omega^2\,, \tag{7.63}$$

where $i = \mathrm{bb}, \mathrm{fb}, \mathrm{ff}, \mathrm{B}, \mathrm{C}$. Because this has the same ω dependence as the ABC flux, it becomes clear that this correction is most important at low energies. At $1\,\mathrm{keV}$ it can reach a few percent while the effect on the overall flux becomes negligible at less than $0.2\,\%$.

The treatment in sections 7.1.1 and 7.1.2 always assumed the Primakoff effect to be elastic. However, when the ion involved is not fully ionised, the Primakoff process can be inelastic by transferring energy to an electron in the shell. The size of this effect is again negligible, which can be seen from ref. [398], where inelastic scattering was investigated in the context of the inverse Primakoff effect. It showed that the inelastic Primakoff cross section is proportional to the incoherent scattering function $S(q)$, whose maximal value is given by the number of bound electrons $z - Q_z$ [399]. The inelastic Primakoff cross section in the plasma is hence at least suppressed by

$$\frac{\sum_z n_z\,(z-Q_z)^2}{n_e + \bar{n}} \lesssim 0.1\,\%\,. \tag{7.64}$$

This suppression can also not be compensated by resonance effects because the plasma frequency, which acts as an effective photon mass, prohibits that any such resonances can take place.

Another process which is not included in the calculation is the electro-Primakoff effect [231]. It is an electron scattering off an ion or electron with an axion emitted from the virtual photon propagator. The total energy emission related to this process was already calculated in ref. [231]. In agreement with their result, it is found to be completely negligible at the order of 4×10^{-5}. To provide a rough estimate on the correction at specific energies, one can allow for a relative scaling between Primakoff and electro-Primakoff of ω^2. This results in a maximal correction between 1 keV and 10 keV of 0.4 %.

Higher-order QED diagrams can also contribute to the Primakoff process. These are at least suppressed by an additional factor of α_{EM}. Also recall that the axion was assumed to be massless even though helioscopes can in principle detect axions of up to ~1 eV. This can lead to corrections which are expected to be smaller than $m_a/\omega \lesssim 0.1\,\%$.

At low energies near the plasma frequency, the dispersion relation in equation (7.5) and hence the Primakoff flux may receive further corrections. Above, the limit of non-relativistic and non-degenerate electrons was taken. However, evaluating the full expression as given in ref. [233] only changes the prefactor of ω_{pl}^2 by 0.3 %. This results in a tiny change of the low-energy cutoff but does not have any significant effect on the flux in the energy range of interest between 1 and 10 keV.

Yet another potential source of uncertainty is the correct screening prescription. As Raffelt pointed out in ref. [231], the time it takes for one plasmon to cross the scattering potential is much smaller than the typical time it takes for an electron to cover the same distance. This is the justification for deriving the effective Primakoff form factor in the static limit, which amounts to first squaring the matrix elements and then averaging over different charge distributions. The two screening prescriptions were also discussed in the context of axion-electron interactions in section 7.1.4. The resulting form factor is included in the Primakoff rates in equations (7.2) and (7.9). In order to quantify the uncertainty associated with this static limit, Raffelt's calculation is generalised to charges moving at constant velocities in appendix E. The relative corrections are found to be of $\mathcal{O}(10^{-4})$, which makes them negligible.

Finally, it has to be mentioned again that the contribution from plasmon conversion is not treated as a correction to the total Primakoff flux. It is also not included in the MC simulation because it depends on the solar magnetic field, which is not tabulated in the solar models. Leaving plasmon conversions aside is justified because the flux component from longitudinal plasmons is only present at energies $\lesssim 0.2$ keV and can therefore be separated from the Primakoff flux with spectral information (see figure 7.4). In contrast, the flux from non-resonant conversions of transverse plasmons is not equally well separated. However, it has a similar dependence on the large scale magnetic field, which means that if it turns out be relevant, this will be marked by pronounced peaks from the resonant conversion of longitudinal plasmons.

ABC flux

The corrections which have already been discussed above and which equally apply to the ABC flux are listed in table 7.3 without further comment. Additional ones are related to approximations in the derivation of the ABC rate in equation (7.39). These include the non-relativistic expansion of the interaction Hamiltonian for the electron, the leading-order multipole expansion of the transition amplitudes and the separation of wave functions in spin and spatial parts [105].

First, the non-relativistic expansion of the matrix elements, which enters into equations (7.31), (7.32) and (7.37), simply discards terms that are suppressed by the electron velocity. An upper bound on the associated corrections can be found by doing a full relativistic calculation of the bremsstrahlung rates (as described in refs. [233, 362]) in the core of the Sun, where relativistic effects are expected to be most significant. The spectral average, is only modified by about $3\,\%$. This can be rescaled to find an estimate for the high-energy tail, where the relativistic corrections are expected to be less than $6\,\%$.

The multipole expansion of the transition amplitudes for both photon and axion emission is given in ref. [105]. Every higher order comes with an additional factor of $\vec{x} \cdot \vec{k}$, the product of the electron's position operator and the axion or photon momentum $\vec{k}$. The size of this factor can be estimated by $\omega/(m_e Z \alpha_{\text{EM}})$, where k is approximated by ω and the distance was taken as the Bohr radius of a hydrogen-like atom. Since this only applies to fb and bb transitions and since higher-order multipoles do not play a role in ff interactions, the size of this correction relative to the overall flux is limited to few percent for most of the energy range.

The separability of the electron wave function into spin and spatial parts is the final simplification entering the derivation of the key equation (7.37). Effectively, this means that spin-orbit interactions are neglected and that an average over the spin states is used instead. However, these interactions distinguish the energy levels by their total angular momentum and the individual fine-structure transitions may not contribute equally to the axion emission. Taking the average over spin-states is therefore only valid if the fine structure is not resolved, limiting the energy resolution of the approach described in section 7.1.4. In the case of the most prominent Fe peak, this corresponds to an energy resolution of $\Delta E \sim 25\,\text{eV}$.[54] In addition, also the total matrix element depends on the transition energy and would receive corrections of the order of $\Delta E/\omega$.

The partial degeneracy of electron in the solar core was only partially considered for the calculation of the ABC flux. While the screening scale is calculated using the full expression in equation (7.13), Pauli blocking is neglected in the rates of equations (7.31) and (7.32). To estimate the effect on the total rate, one can evaluate the full ion-bremsstrahlung phase space integrals with and without Pauli blocking around the solar core. The resulting suppression factor can be used for all radii. In the case of electron-electron bremsstrahlung, it is squared due to the two electrons in the final state. Since degeneracy effects are expected to be strongest at the core, this estimate should be very conservative. It comes out with a maximal suppression of the total ABC flux by $7\,\%$. In principle, it would be possible to replace equations (7.31) and (7.32) by full expressions including Pauli blocking. However, since the ABC flux is affected by several other large uncertainties and due to the computational cost

---

[54] This estimate corresponds to the spread between fine-structure states of the relevant Lyman-alpha transition and is taken from the NIST database [397].

of such high-dimensional integrals, the full phase-space integrals were not included in this work. For a high-precision measurement of the ABC flux, this calculation should nonetheless be done in the future.

Besides the right evaluation of the screening scale including degeneracy effects, it is important to apply the right screening prescription, i.e. the right modified photon propagator inside the plasma. As already discussed in section 7.1.4, this leads to additional uncertainties of the order of at least a few percent. In the case of the Primakoff flux it was possible to estimate the next-to-leading order effects in appendix E.

Bremsstrahlung rates in equations (7.31) and (7.32) are computed in the Born approximation. This means that incident electrons are described by plain waves. However, long-range effective forces can significantly modify the wave solution of free non-relativistic particles and thereby enhance the scattering amplitude. The process is an example of Sommerfeld enhancement. For an unscreened Coulomb potential, the size of this effect can be well approximated by including the Elwert factor [361, 400] in the phase space integral. Including the Elwert factor in equations (7.31) and (7.32) enlarges the results by approximately 20 %. In the solar plasma, where charges are screened, the effect will be much weaker. One can use the estimate provided in ref. [361] and conclude that the individual rates (7.31) and (7.32) should underestimate the true value by less than 10 %. The effect on the total ABC flux will be smaller again.

Finally, a similar effect is also hidden in the OP opacities. As Redondo pointed out [105], they also used Coulomb wave function in their calculations. As above, this means that the effect of enhanced wave functions towards the centre of the scattering potential is overestimated since charge screening is not taken into account for the incident waves. The approximation of Coulomb waves should, however, become more precise with larger charges of the target nuclei because Coulomb wave functions and plain waves only differ at radii given by the size of electronic orbitals around the nucleus. This is yet another reason to use the analytical results in equations (7.31) and (7.32) including the Born approximation for the scattering of light nuclei or two electrons and to only apply equation (7.39) with the OP data for heavier elements.

To conclude the discussion of incident electron waves, note that there are two corrections of the calculation in section 7.1.4 – one increases and one decreases the total ABC rate. The true Bremsstrahlung rates should be larger due to Sommerfeld enhancement. On the other hand, the Coulomb wave functions employed by the OP overestimate the rates of electrons scattering off heavier nuclei but to a smaller degree. A precise quantification of the combined effect is difficult but a very conservative estimate would be a systematic uncertainty on the ABC flux of less than 10 %.

Interference effects

So far, the Primakoff and ABC fluxes were treated as completely independent of one another. There is however one possible effect which can invalidate this approach – interference. It can only occur between two production processes with identical initial and final states. Looking at the diagrams in figure 7.1, it becomes clear that

Model	N_a^{P}		N_a^{ABC}	
	Mean	Error	Mean	Error
AGSS09	1	1.2 %	1	1.4 %
GS98	1.052	1.3 %	1.051	1.6 %

Table 7.4. Statistical uncertainties ("Error") of the integrated Primakoff and ABC fluxes calculated with the MC simulation. The mean values are normalized to the mean of the AGSS09 model. Note that the Primakoff flux is subject to additional systematic uncertainties of $\lesssim$1 %. The ABC calculation is expected to be significantly less accurate with a precision of ~10 % (see table 7.3).

only Primakoff and Compton processes can interfere. As pointed out in ref. [231], this interference is suppressed by ω/m_e because in the non-relativistic limit the two final state electrons have opposing spin. Therefore, the interference can only give a significant contribution to the Primakoff flux when g_{ae} is very large, compensating for the ω/m_e suppression and the fact that the Compton process only contributes a small fraction to the ABC flux. Such a large electron coupling would necessarily result in a dominant ABC flux. Precision measurements like the one proposed in section 7.4 require the Primakoff flux to dominate, so that this interference can safely be neglected.

Going beyond the processes explicitly included in the calculation, there could also be an interference between electro-Primakoff and bremsstrahlung. However, the same suppression by ω/m_e is found by the same argument as above. Taking into account that the electro-Primakoff was already negligible on its own, one can conclude that this interference term is also insignificant.

Sun-Earth distance and positioning

A final effect which is not included in the computation of axion fluxes is the dependence on the varying distance between Sun and Earth d_{E}. Over the course of a year, d_{E} oscillates around its average by about 1.7 % [381]. Consequently, this effect has to be taken into account when a sizeable axion signal is detected. Since d_{E} can, in principle, be measured with high precision, this does not need to be treated as an additional source of uncertainty. All numerical values given in this thesis can also be rescaled to the correct mean value for any given experiment.

A similar, biannual effect applies to the relative positioning of Earth with respect to the Sun's equator. This affects the plasmon flux because the large-scale solar magnetic fields are not spherically symmetric. The resulting geometrical factor may be determined for the data taking times of an experiment and then be used to rescale the equations presented in this work.

7.3.5 Summary and consequences

The extensive review of the solar axion flux above had two objectives. First, the calculation was updated to include all recent advances in solar modelling, opacity evaluations and calculations of axion emission rates. Second, the corresponding uncertainties were quantified. For two different solar models, the uncertainty coming from the variation of input parameters within their error bars was translated to the axion flux using a Monte Carlo simulation. The resulting uncertainties of the total integrated axion flux $\mathcal{N}_a$ are summarised in table 7.4. For each solar model, the uncertainty is only at the percent level. However, for both Primakoff and ABC flux, the resulting counts are systematically higher by about 5 % for photospheric, GS98-type models. This does not invalidate the results, but it is a consequence of the solar metallicity problem.

Additional, systematic uncertainties arise from approximations made in the calculation of the rates. One would naively expect that the associated corrections would be of similar size for either type of solar model. However, some of them will also have a slightly larger impact on high-Z models than on low-Z ones, which could in principle change the relative factor between the two models considered. The relative size of the uncertainties was estimated in section 7.3 and they are listed in table 7.3. While the Primakoff flux is only affected by corrections $\lesssim 1\,\%$, the ABC flux could by modified by $\sim 10\,\%$. It is therefore clear that only the Primakoff flux currently allows high-precision measurements.

To give a concrete example, one can study how the updated calculations affect solar axion bounds. To this end, the latest data and limits from the CAST experiment [208] are used. Figure 7.11 shows the profile likelihood ratios resulting from this measurement using the updated axion flux for different solar models. Clearly, the overall effect of the systematic differences on the limits from different models is rather small and the intrinsic uncertainties for every model (not shown) are smaller still. This is not unexpected, as the uncertainties only enter as the fourth root in the effective coupling parameter.

In conclusion, the quantification of uncertainties in this thesis shows that the solar axion flux is known to a high enough precision to put robust bounds on axion models. Furthermore, the Primakoff flux may allow precision measurements once an axion signal is detected. In this case, controlling the uncertainties is crucial in order to measure properties of the axion [195, 196] or the Sun [3, 346, 347]. An example for how both could be achieved simultaneously is given in the next section.

7.4 Disentangling axion and solar models

The discussion of uncertainties showed that the intrinsic uncertainty of each model is typically much smaller than the difference between high-Z and low-Z solar models (see table 7.4). This means that it is in principle possible to distinguish the two limiting cases. The problem with this naive observation is a degeneracy between the solar metallicity and the axion coupling constants. Smaller Z values result in a smaller axion flux, which may always be compensated by a larger value of the unknown coupling constant. Nonetheless, a crucial property of the axion-photon coupling allows to circumvent this degeneracy. As it was discussed in chapter 5, the phenomenological coupling constant $g_{a\gamma\gamma}$ is generated from an anomaly and the associated anomaly coefficient takes on discrete values for different viable axion

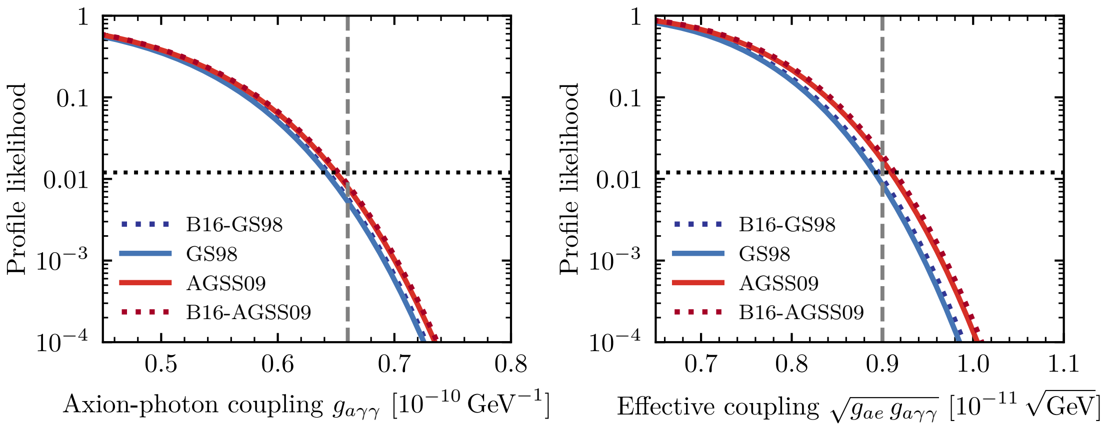

Figure 7.11. Profile likelihoods of the axion-photon coupling and the combination $\sqrt{g_{ae}\, g_{a\gamma\gamma}}$. The left and right panel assume that the flux is dominated by Primakoff or ABC axions, respectively. The latest CAST data from 2017 [208] and the implementation in ref. [342] are used. The different coloured lines represent a selection of solar models. 95 % confidence limits are indicated by dashed grey lines and the corresponding nominal likelihood thresholds by dotted black lines. Note that the CAST 2017 limit on $g_{a\gamma\gamma}$ [208] uses the BP04 model and that the reference limit for $\sqrt{g_{ae}\, g_{a\gamma\gamma}}$ was only derived for older data by the CAST collaboration [352].

models ($g_{a\gamma\gamma} \propto E/N + \text{const.}$). This means that also $g_{a\gamma\gamma}$ is quantised once the mass of the axion and thereby f_a is known (see equation (2.28)). Hence, the described degeneracy is between one parameter – the metallicity Z – which is constrained to fall within a relatively narrow interval and another parameter – the coupling $g_{a\gamma\gamma}$ – which can only take discrete values. Combining the knowledge about these two parameters can allow to measure them separately.

A comprehensive list of values of $g_{a\gamma\gamma}$ in hadronic axion models with a single additional coloured fermion is given in refs. [108, 115]. The authors consider models to be *preferred* if the additional heavy quark is sufficiently unstable and if the model has no Landau pole below the Planck scale. By considering hadronic axions with the above properties, which is sometimes called the *axion window*, one can break the degeneracy between the solar metallicity and the coupling constant. However, it is important to note that a specific E/N value does in general not correspond to a unique UV model. For instance, among the models that will be considered below, there are multiple ones with $E/N = 2/3$ and $E/N = 8/3$. The number of possible E/N values could also be drastically increased if one would extend the models under consideration to KSVZ axion models with more than one new heavy fermion (see e.g. ref. [124] for a recent review on the landscape of axion models). It would therefore require additional information to distinguish these. Nevertheless, even a rough measurement of E/N would be a severe restriction on possible UV models and when one goes beyond typical values for E/N, cosmological problems or an extended UV sector commonly appear.

Having established a set of benchmark hadronic axion models, it is interesting to investigate whether the most sensitive solar axion experiment IAXO can distinguish the different models within the axion window. Moreover, could it even find a preferred combination of a solar and an axion model? These questions can only be answered by taking all of the relevant uncertainties into account, which include the ones from the MC simulation and the systematic ones discussed in section 7.3 as well as the statistical uncertainty of the measurement itself.

Since hadronic axion models have a negligible axion-electron coupling, the corresponding solar axion spectrum will be Primakoff-dominated, $\gamma \gtrsim 10^7$.[55] Because the relative uncertainties of the Primakoff flux in the left panels of figure 7.7 and 7.9 are approximately constant for all relevant energies[56], the solar uncertainty can be modelled by a simple re-scaling of the axion spectrum. This is done by introducing a scaling constant C such that the overall signal is proportional to $(C g_{a\gamma\gamma})^4$. C is normalized to unity for the B16-AGSS09 model, which is the latest low-Z solar model. With these definitions, IAXO does not directly constrain $g_{a\gamma\gamma}$ but instead $C g_{a\gamma\gamma}$ – a combination of axion and solar properties.

It is illustrative to make a plot of IAXO's predicted sensitivity in the plane spanned by m_a and $C g_{a\gamma\gamma}$. Because C is different for high- and low-Z solar models, each combination of a solar model with a hadronic axion model can be plotted separately. This results in low-Z (red) and high-Z (blue) bands in figure 7.12, which are only clearly visible as individual bands in the magnifications (right panel). Overlapping bands appear in a darker, purple shading.

The underlying uncertainties in the axion model parameters are the reason why each combination of a solar and an axion model defines a narrow band rather than a line. The width of each band includes the statistical uncertainty of the solar model in addition to the QCD uncertainties of m_a and $g_{a\gamma\gamma}$, which are given in equations (2.28) and (3.31). The uncertainty from $\widetilde{C}_{a\gamma\gamma}$ dominates the total theoretical uncertainty of axion models when E/N is close to $\widetilde{C}_{a\gamma\gamma}$, i.e. for models with small couplings to photons. This effect is clearly visible in figure 7.12, where the bands get wider with smaller couplings.

The red, blue and purple bands indicate the theoretical uncertainty from QCD and the axion flux. In addition, experimental uncertainties are shown as grey regions. To estimate these, a likelihood ratio analysis was performed using an Asimov data set to find the expected exclusion contours at the 95 % confidence limit (CL). Similar studies have been performed in refs. [195, 196]. The same method is applied for seven benchmark values of m_a and $g_{a\gamma\gamma}$, using the baseline IAXO parameters as given in ref. [193]. The points are chosen by hand to illustrate qualitatively-different potential experimental outcomes. They are marked with black crosses in figure 7.12 and are surrounded by their respective expected 95 % CL contour.[57]

The two points at the bottom with the smallest coupling constants demonstrate that to locate the axion within a specific model band necessarily requires a mass

---

[55] Following the arguments in section 7.1.3, the flux generated by plasmon conversions in large scale magnetic fields is considered as either small or separately detectable.

[56] While not exactly true, this is an acceptable approximation for estimating IAXO's potential. In case of a detection, the data analysis (in the m_a-$g_{a\gamma\gamma}$ plane) would need to be repeated with various solar models (not by a simple rescaling) in order to make statements about the nature of the signal.

[57] For simplicity, only the vacuum setup is simulated. If the conversion volume is filled with a buffer gas, the experimental sensitivity would improve for higher axion masses [193, 194].

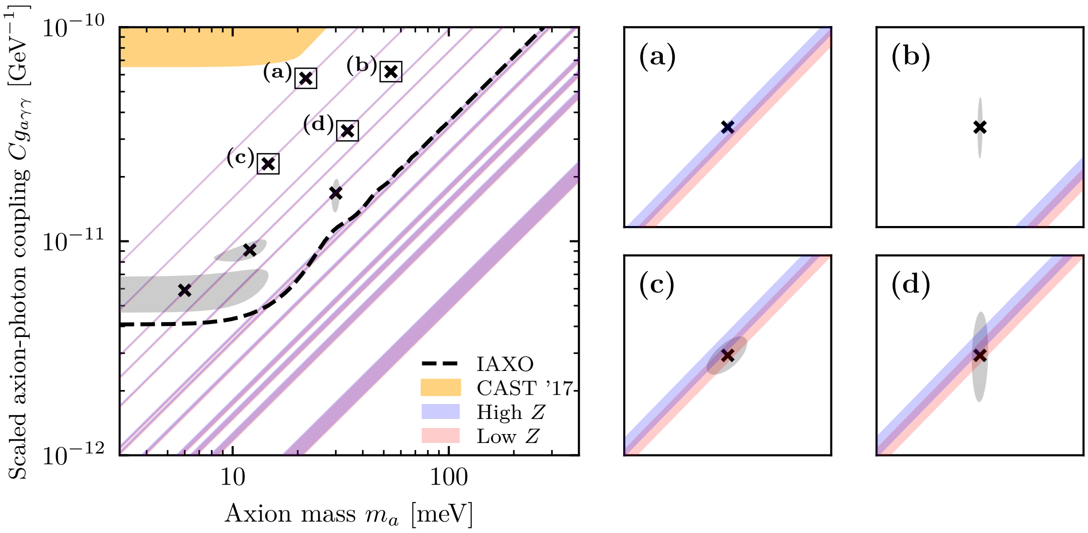

Figure 7.12. IAXO sensitivity and benchmark cases in the m_a-$Cg_{a\gamma\gamma}$ parameter plane. Included are all hadronic axion models in the *axion window* for high- and low-Z solar models (blue and red shading, respectively; their overlaps appear in a darker, purple shading). *Left:* Black crosses mark seven benchmark points for which the expected 95 % confidence limits (grey-shaded areas) of a detection with the baseline IAXO setup [193] are determined. The black line marks the expected IAXO sensitivity, while the yellow region has already been excluded by CAST [208]. *Right:* Magnifications of the most precise measurements, labelled (a) to (d).

measurement. For IAXO this is only possible if the axions mass is at least a few meV and the coupling is sufficiently large [195, 196]. A detailed discussion and the full region of parameter space where a mass measurement is possible is given in ref. [195].

In all cases with $g_{a\gamma\gamma} \gtrsim 10^{-11}\,\mathrm{GeV}^{-1}$ and $m_a \gtrsim 10\,\mathrm{meV}$, the axion parameters can be located on a specific axion band, which means that there is only one possible value of E/N within the *axion window* that would explain the observation. This is of course only true if the observed particle actually falls on such a band. The cross labelled (b) is an example for a situation where this would not be the case and the detected axion would not be included in the list of preferred models.

At large coupling values, the experimental precision becomes so good that the expected 95 % CL contour is only visible in the corresponding magnifications in the right part of figure 7.12. It is also only at this accuracy that the high- (blue) and low-metallicity (red) solar model bands are clearly distinguishable. The two bands overlap due to the combination of solar model and QCD uncertainties, of which the latter dominate the overall uncertainty. In fact, without QCD uncertainties, the two bands would not overlap.

It is even possible to obtain a hint from IAXO towards higher or lower metallicity in some serendipitous cases. For instance, a measurement as indicated in subpanel (a) of figure 7.12 could only be explained by the combination of a specific value of E/N and a high-Z solar model. Example (c) could also be interpreted as a hint towards lower Z even though the picture is less clear in this case. In contrast, the experimental uncertainty smears out any potential information on the solar metallicity in example (d).

In conclusion, IAXO can indeed differentiate between different hadronic axion models in the *axion window*. This is approximately possible for all parameters in the mass-coupling plane, for which IAXO can determine the (non-zero) axion mass, as described in refs. [195, 196]. In the case of a strong signal in IAXO, it may also be possible to infer a combination of a solar model and axion model at the same time. Indeed for such a fortuitous case and making the assumption of a reasonably simple KSVZ model, it is possible to distinguish low-Z and high-Z models without making use of the g_{ae} coupling (as it was done in ref. [346]). It is, however, also possible that a detected particle is not in the *axion window*. This could for instance be an axion model which includes multiple heavy coloured fermions in various representations of the SM gauge groups [108]. In such a case, the degeneracy between metallicity and the axion coupling would still persist but by constraining the value of E/N including the solar model uncertainty, a large set of UV axion models could be ruled out nonetheless.

7.5 IAXO and the 14.4 keV line

Nuclear interactions of axions offer an additional opportunity for axion detection and model distinction. As discussed in section 7.1.5, axions emitted by thermally excited nuclei could be detected by helioscopes, because they fall into the right energy range. Out of all possible magnetic dipole transitions, the one of ^{57}Fe generates by far the strongest solar axion flux (see appendix F), which is why only the corresponding 14.4 keV line is considered in the following. The search for these ^{57}Fe axions has a long history. The currently most powerful helioscope, the CERN Axion Solar Telescope (CAST) [203], as well as CUORE [378] and, more recently, XENON1T [212], have searched for axions produced in this transition and provided constraints on the axion-nucleon coupling. With the new helioscope IAXO under construction, it is again necessary to asses its potential to detect axions from this nuclear transition. This is because detecting axions at 14.4 keV comes with its own unique experimental challenges and it is helpful to provide estimates of efficiencies for a variety of different setups.

In section 7.1.5, the power of the ^{57}Fe line as a function of g_{aN}^{eff} was updated to include the latest new theoretical developments like revised nuclear matrix elements and the latest solar models. As a next step, this section investigates how the flux depends on the axion model and which models can be used as benchmarks in the following sensitivity studies. Finally, various experimental setups of IAXO are considered in order to provide a guide for the techniques required to maximise the efficiency to detect axions from ^{57}Fe.

7.5.1 Benchmark models

The axion-nucleon couplings which enter the relevant effective coupling g_{aN}^{eff} can be expressed in terms of the dimensionless axion-quark coefficients c_q^0 [54]. Using this, equation (7.43) can be written as[58]

$$\frac{\Gamma_a}{\Gamma_\gamma} = 5.81 \times 10^{-16}(1 + 3.28\, c_u^0 - 9.97\, c_d^0 + 0.52\, c_s^0 + 0.16\, c_c^0 + 0.12\, c_b^0 + 0.048\, c_t^0)^2 \left(\frac{m_a}{1\,\mathrm{eV}}\right)^2 . \tag{7.65}$$

The first typical benchmark model is the KSVZ [109, 110] model with $c_q^0 = 0$. This results in

$$\left.\frac{\Gamma_a}{\Gamma_\gamma}\right|_{\mathrm{KSVZ}} = 5.81 \times 10^{-16} \left(\frac{m_a}{1\,\mathrm{eV}}\right)^2 . \tag{7.66}$$

In the DFSZ model [116, 117], $c_{u,c,t}^0 = \cos^2(\beta)/3$ and $c_{d,s,b}^0 = \sin^2(\beta)/3$ corresponding to

$$\left.\frac{\Gamma_a}{\Gamma_\gamma}\right|_{\mathrm{DFSZ}} = 5.81 \times 10^{-16} \times (1 + 1.16 \cos^2\beta - 3.11 \sin^2\beta)^2 \left(\frac{m_a}{1\,\mathrm{eV}}\right)^2 . \tag{7.67}$$

One can see that the axion emission rate is accidentally suppressed in KSVZ models due to a small neutron coupling. In DFSZ models, on the other hand, it can get enhanced by up to a factor of ~ 4 with respect to KSVZ.

Equation (7.65) shows that for downphilic axions, i.e. axions with large $c_d^0 \gg c_u^0$, the emission rate is enhanced if the model-independent term normalized to one in (7.65) is not the dominating contribution. Such enhancements naturally occur in non-universal DFSZ models. For instance, the M1 model of ref. [121] has quark couplings $c_u^0 = c_c^0 = \sin^2\beta$, $c_t^0 = -\cos^2\beta$, $c_d^0 = c_s^0 = \cos^2\beta$ and $c_b^0 = -\sin^2\beta$, leading to

$$\left.\frac{\Gamma_a}{\Gamma_\gamma}\right|_{\mathrm{M1}} = 5.81 \times 10^{-16} \times (1 + 3.32 \sin^2\beta - 9.50 \cos^2\beta)^2 \left(\frac{m_a}{1\,\mathrm{eV}}\right)^2 . \tag{7.68}$$

At small β, this can give an enhancement of $\mathcal{O}(60)$ with respect to the KSVZ model. Other non-universal DFSZ models feature similar enhancements of g_{aN}^{eff} but they can have different values for the axion-photon coupling. This is important since the detection of axions in helioscopes always uses this coupling. In particular, the non-universal model $\mathcal{T}_2^{(u)}$ of ref. [122] features the largest axion coupling to photons among the general class of non-universal DFSZ models with two Higgs doublets (see table 5 in [123]).

7.5.2 Astrophysical bounds

The strongest bounds on axion-nucleon interactions today come from astrophysical observations. Because the Sun's interior dynamics are particularly well understood, one can derive a solid bound on g_{aN}^{eff} from a solar energy loss argument. Using the

---

[58] In deriving this equation, the standard relation between m_a and f_a (as in equation (2.28)) was used.

revised axion rate in equation (7.43) and the solar model B16-AGSS09, the total axion luminosity (energy loss rate) via the ^{57}Fe transition comes out as

$$L_a = 8.38 \times 10^9 (g_{aN}^{\text{eff}})^2 \, L_\odot \ . \tag{7.69}$$

Exotic solar energy loss was constrained in ref. [401] to be smaller than 3 % of the total solar luminosity $L_\odot$. This can be translated to a bound on the effective nucleon coupling,

$$\left| g_{aN}^{\text{eff}} \right| \leq 1.89 \times 10^{-6} \ . \tag{7.70}$$

Notably, this bound is about a factor of two more stringent than the previous constraint, $g_{aN}^{\text{eff}} \leq 3.6 \times 10^{-6}$ [203]. Besides the enhanced emission rate [377] and the updated solar model [379], this is also due to ref. [203] having excluded only $L_a > 0.1 \ L_\odot$.

While the solar bound is very robust, it is not the strongest one derived from astrophysics. This role falls to constraints which use the neutrino burst observed in coincidence with the SN1987A event [263, 265, 402–408]. The most recent analysis [265] derived a bound on a combination of axion-proton and axion-neutron couplings,

$$g_{an}^2 + 0.6 \, g_{ap}^2 + 0.5 \, g_{an} \, g_{ap} \lesssim 8.3 \times 10^{-19} \ . \tag{7.71}$$

Because this is not the combination relevant for solar axions from the ^{57}Fe transition, it has to be translated into a bound on g_{aN}^{eff}. This is done by choosing the ratio between the proton and neutron couplings such that the left-hand side of this equation is minimal, $g_{ap}/g_{an} \approx -0.2$, while keeping g_{aN}^{eff} constant,

$$\left| g_{aN}^{\text{eff}} \right| \lesssim 1.1 \times 10^{-9} \ . \tag{7.72}$$

This strong constraint can only be overcome by highly advanced helioscope setups as will become clear in the next sections. Nonetheless, the observations on which this bound is based are affected by their own substantial uncertainties (e.g. relying on a single supernova event). It would therefore be comforting to have independent confirmation in more controlled setups.

Finally, axion-nucleon interactions can also be constrained from x-ray observations of various NS [409–414]. These bounds are, however, subject to several uncertainties and do not always agree with each other, not even when referring to the same star [123]. In any case, all these analyses suggest a limit of $\sim 10^{-9}$ on some combination of axion-nucleon couplings.[59]

7.5.3 Experimental configurations

Before going into specific experimental aspects, one can make a few general considerations that can provide help with the experimental design. To detect the ^{57}Fe line, the signal to noise ratio in the single relevant energy bin, which contains all of

---

[59] In most of these analyses, the limit applies only to the axion-neutron coupling.

the signal events, should be maximised. There are two main contributions to the background for the measurement of the nucleon line. One is the usual background rate of the detector itself (e.g. cosmic rays, environmental gammas, intrinsic detector radioactivity), which usually grows linearly with the area of the detector and is typically proportional to the spectral size of the signal bin. This means that, in the case of the expected narrow signals, it can be reduced by making use of a detector with good energy resolution. Second, there is the physics background due to Primakoff production. This grows with the photon coupling. Importantly, as the Primakoff spectrum is continuous, it also grows linearly with worsening energy resolution. Combining these two effects leads to the following figure of merit,

$$f \propto \frac{S}{\sqrt{B}} \propto \frac{\epsilon_o \epsilon_d \; g_{a\gamma\gamma}^2}{\sqrt{\Delta E_d}\sqrt{ba + g_{a\gamma\gamma}^4 \kappa \epsilon_o \epsilon_d}} \,. \tag{7.73}$$

In this equation, $\epsilon_{o,d}$ are the optics and detector efficiencies, ΔE_d is the energy resolution (FWHM) of the detector, b is the (spectral) background rate per area and a is the signal spot area on the detector. κ quantifies the Primakoff flux in the ^{57}Fe signal bin and is implicitly defined in equation (7.77).

The figure of merit clearly identifies the parameters, which need to be optimised. While good energy resolution is critical, this should not be offset by too large a background. Similarly, with focusing optics, the detector area and therefore the background contribution can be reduced . However, there is a balance because such x-ray optics may have an efficiency significantly smaller than one. All of these trade-offs will become more apparent when considering explicit setups below.

BabyIAXO

BabyIAXO [194, 415] is an intermediate stage of the full IAXO experiment. It serves as an experimental pathfinder but is already expected to substantially advance the exploration of the axion parameter space. It will be able to study QCD axion models and investigate stellar cooling hints along with other well motivated sections of the parameter space [124, 193]. The BabyIAXO experiment is mainly designed to measure Primakoff axions in the energy range from 1 keV to 10 keV with the peak of the solar axion flux spectrum at ∼3 keV. BabyIAXO will consist of two magnetic bores of 10 m length and 70 cm diameter, each with an average magnetic field strength of about 2 T. Together with newly developed x-ray optics and detector systems providing higher energy resolution and lower background, BabyIAXO will be the first helioscope to exceed the sensitivity of the CAST experiment. The magnet bores of BabyIAXO are of similar diameter as those of IAXO (60 cm) and IAXO+ (80 cm), thus the experience from the optics and detector development for BabyIAXO can later be applied to the upgraded configurations [194]. The first detection system for BabyIAXO is chosen to be microbulk Micromegas technology. These detectors have proven background levels as low as $10^{-7}\,\mathrm{keV}^{-1}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ [416] and a high detection efficiency for Primakoff photons below 10 keV. A variety of other detector types, like silicon drift detectors (SDD), metallic magnetic calorimeters (MMC) and transition edge sensors (TES), are also studied for BabyIAXO aiming to optimise the energy resolution for precision measurements of the axion spectrum [194].

	BabyIAXO			
	baseline	no optics	optimised optics	high energy resolution
	BabyIAXO$_0$	BabyIAXO$_1$	BabyIAXO$_2$	BabyIAXO$_3$
B [T]	2	2	2	2
L [m]	10	10	10	10
A [m^2]	0.77	0.38	0.38	0.38
t [year]	0.75	0.75	0.75	0.75
b $\left[\frac{1}{\mathrm{keV\,cm^2\,s}}\right]$	10^{-7}	10^{-6}	10^{-7}	10^{-5}
ϵ_d	0.15	0.9	0.5	0.99
ϵ_0	0.013	1	0.3	0.3
a [cm^2]	0.6	3800	0.3	0.3
$r_\omega = \frac{\Delta E_d}{14.4\,\mathrm{keV}}$	0.12	0.12	0.12	0.02

	IAXO		**IAXO+**	
	low background	high energy resolution	low background	high energy resolution
	IAXO$_\mathrm{b}$	IAXO$_\mathrm{r}$	IAXO$^+_\mathrm{b}$	IAXO$^+_\mathrm{r}$
B [T]	2.5	2.5	3.5	3.5
L [m]	20	20	22	22
A [m^2]	2.3	2.3	3.9	3.9
t [year]	1.5	1.5	2.5	2.5
b $\left[\frac{1}{\mathrm{keV\,cm^2\,s}}\right]$	10^{-8}	10^{-6}	10^{-9}	10^{-6}
ϵ_d	0.99	0.99	0.99	0.99
ϵ_0	0.3	0.3	0.3	0.3
a [cm^2]	1.2	1.2	1.2	1.2
$r_\omega = \frac{\Delta E_d}{14.4\,\mathrm{keV}}$	0.02	$\frac{5}{14400}$	0.02	$\frac{5}{14400}$

Table 7.5. List of experimental parameters adopted for all helioscope configurations which are considered in figures 7.13 and 7.14. B is the magnetic field of the helioscope, L its length and A the cross-sectional area. t is the total observation time in which the helioscope is pointed at the Sun. All of these values are taken from ref. [193]. $\epsilon_{o,d}$ are the efficiencies of the optics and detector, b is the spectral background rate per detector area and r_ω is the relative spectral resolution of the detector. The setup BabyIAXO$_0$ is the baseline BabyIAXO, BabyIAXO$_1$ is a version without optics, BabyIAXO$_{2,3}$ assume optics optimised for the 14.4 keV line with BabyIAXO$_3$ including also a good energy resolution. For each of the more advanced setups IAXO and IAXO+, a version with minimal background and one with optimised energy resolution is considered.

The baseline BabyIAXO configuration is not optimal for the measurement of ^{57}Fe axions. For instance, simulations of the Micromegas detector system predict an ionization probability at a conversion length of 3 cm at the given argon gas mixture and pressure of only around 15 %. Also the BabyIAXO optics are not optimised for 14.4 keV photons, where the expected optics efficiency is merely 1.3 % (see BabyIAXO$_0$ configuration in table 7.5). For these reasons, some possible modification of the BabyIAXO experiment are considered in order to boost its sensitivity to axions from ^{57}Fe.

The sensitivity of the Micromegas detector to 14.4 keV photons can for instance be enhanced by changing the gas mixture and adjusting the pressure. Gas mixtures with an inert gas of higher atomic number, like xenon, show a higher conversion efficiency for 14.4 keV photons, compared to argon. With a suitable conversion length, an efficiency of >90 % should be reachable at atmospheric pressure.

Regarding the optics, one could consider the radical approach of removing them all together. Currently, microbulk Micromegas [417] are developed to have a square active area of 25 cm × 25 cm [418], which means that the whole bore opening can be covered with a few detector tiles. This could improve the detection efficiency, because losses from inefficient x-ray optics are avoided. However, the detector will have a higher background due to the larger conversion volume. This *no optics* configuration is denoted as BabyIAXO$_1$ in table 7.5 and in the figures below. As BabyIAXO features two bores, the optics could also only be removed from one of the two bores, while the other one could be operated in parallel with optics and a smaller detector.

An alternative detection concept to Micromegas are silicon drift detectors (SDD). Such a detector consists of a thick negative doped layer, which is fully depleted by a negative bias voltage and positive doped contacts and strips on both sides of the layer. The incoming x-ray radiation generates electrons in the depleted zone that are then drifted towards the anode at the end of the layer. This detector type is already considered as an addition for the baseline measurements at BabyIAXO [194] and it can reach a sensitivity of up to 50 % to 15 keV photons [419]. Because SDDs are only produced in small pixels and a maximal detector size of a few millimetres, x-ray optics are mandatory in this detection scheme.

In addition, the optics could be optimised for focussing high-energy x-rays using multilayer coating techniques [420]. Such an optical system was already employed by the Nuclear Spectroscopic Telescope Array (NuSTAR) [421] to detect photons in the range of 5 keV to 80 keV. A BabyIAXO system, which does not require imaging optics, could reach an even higher throughput. A realistic estimate in this scenario would be $\epsilon_0 = 0.3$, which is the value adopted in the analysis below. The corresponding setup is denoted BabyIAXO$_2$ in table 7.5.

Another detector suited to the ^{57}Fe line could be one based on Cadmium-Zinc-Telluride (CZT) semiconductors. While the detection principle is similar to the one of SDDs, CZT provides a higher ionization probability for the photons such that only 300 µm of CZT have an ionization probability of >99 % at 14.4 keV. CZT detectors are already used in experiments focusing on the detection of hard x-rays like NuSTAR [422]. These detectors can reach an energy resolution of 2 % in the relevant range [423], which makes them powerful in discriminating the peaked signal from the continuous Primakoff background. However, the good energy resolution comes at the price of an increased background compared to SDDs (see table 7.5). Furthermore, it is currently only possible to produce detectors with an active area of up to 2 cm × 2 cm [424], which again makes an efficient x-ray optic necessary for the use of CZT detectors at BabyIAXO. The corresponding setup, which focusses on energy resolution, is denoted BabyIAXO$_3$.

All of these modifications are assumed to be implemented in only one of the two BabyIAXO bores. This is why the cross-sectional area A in all of the optimised setups in table 7.5 is smaller than the baseline value of BabyIAXO$_0$.

IAXO and IAXO+

Construction of the full scale IAXO experiment is expected to begin during the BabyIAXO data-taking period. IAXO+ is an even more ambitious experimental configuration with a significantly larger magnet as well as improved optics and detector designs [123, 193].

In this thesis, two types of setups are considered for both IAXO and IAXO+, which are summarised in table 7.5. The first is labelled by an index b and assumes optimistically that significant improvements in the detector backgrounds can be achieved. At the same time, the energy resolution is only set to a moderate value of 2 % at 14.4 keV, which is a very realistic figure for CZT detectors as discussed above. This setup is therefore aimed at reducing the intrinsic detector background at the expense of a larger physics background from Primakoff axions. The contrary option would be to optimise the energy resolution, thereby reducing the Primakoff background, at the expense of a larger detector background. These configurations are labelled with an index r. They assume a resolution at the level of the Doppler broadening of the ^{57}Fe line. Such a precise energy resolution is realistic when new technologies like microcalorimeters operated at mK temperatures are employed [425, 426]. Metallic magnetic calorimeters (MMCs) are also studied in the IAXO collaboration for precision measurements of the axion spectrum in the energy range $<$ 10 keV. First measurements have shown an energy resolution of 6.1 eV (FWHM) at the 5.9 keV ^{55}Fe peak [425]. Because sufficiently low background still needs to be established, a somewhat larger background rate is assumed in these high-resolution setups in table 7.5.

The IAXO+ setups are anticipated to include further improvements of the detection system combined with an enhanced magnetic field, cross-sectional area, length etc. as outlined in ref. [193].

7.5.4 Sensitivity evaluations

The sensitivity to g_{aN}^{eff} of a helioscope depends in general on several other axion parameters: on $g_{a\gamma\gamma}$ and m_a through the conversion probability $P_{a\to\gamma}$ (see equation (4.3)) and also on $g_{a\gamma\gamma}$ and g_{ae} because of the Primakoff and ABC contributions to the background. In the following, the ABC background is neglected because it can in principle be completely absent (unlike the Primakoff flux) and because it decays more quickly at larger energies and is therefore subdominant to the Primakoff background in most of the common axion models. To further simplify the discussion, the axion can be assumed to be massless or very light. Nonetheless, effects of finite axion masses are considered separately below.

Massless axions

If the axion mass is assumed to be small ($\lesssim 20\,\mathrm{meV}$), the conversion probability of an axion to a photon $P_{a\to\gamma}$ becomes energy independent, as shown in equation (4.4). The helioscope signal can be calculated by integrating the product of the spectral axion flux over the range of given energy bins. Because the FWHM of the Doppler-broadened iron peak is only $\sim 5\,\mathrm{eV}$, it can be assumed that the whole signal is always in one energy bin.[60] This means that the expected number of signal events μ_signal can be calculated from the total ^{57}Fe flux Φ_a^Fe57 as

$$\begin{aligned}\mu_\mathrm{signal} &= \Phi_a^\mathrm{Fe57}\ P_{a\to\gamma}\ A\ t\ \epsilon_o\ \epsilon_d && (7.74)\\ &\propto \left(g_{a\gamma\gamma} g_{aN}^\mathrm{eff}\right)^2\ . && (7.75)\end{aligned}$$

The background consists of two contributions. First, the detector background, which is quantified by the background level b and which can be measured accurately at times when the magnet bores are not pointed at the Sun. Second, the tail of the Primakoff spectrum, which may act as an additional background. The expected background events μ_back are therefore given by

$$\begin{aligned}\mu_\mathrm{back} &= \int_{E^*-\frac{\Delta E_d}{2}}^{E^*+\frac{\Delta E_d}{2}} \left(\frac{\mathrm{d}\Phi_a^\mathrm{P}}{\mathrm{d}\omega}\ \epsilon_o\epsilon_d\right)\mathrm{d}\omega\ P_{a\to\gamma}At + bat\Delta E_d && (7.76)\\ &\simeq \left(g_{a\gamma\gamma}^4 \kappa\epsilon_o\epsilon_d + ba\right)\Delta E_d t\ . && (7.77)\end{aligned}$$

The Primakoff flux Φ_a^P as well as the efficiencies ϵ_o and ϵ_d are in general functions of ω. In case of a sufficiently small energy resolution ΔE_d, one can average over the energy and describe the Primakoff background using the constant κ. If the Primakoff background to the ^{57}Fe-peak is detectable at 14.4 keV, it is clear that there is a much stronger Primakoff signal at smaller energies and $g_{a\gamma\gamma}$ has been measured very precisely. Therefore, the expected contribution from Primakoff axions to the number of background events μ_back is either known or it is negligible compared to the intrinsic detector background.

The p-value of the Poisson-distributed observed number of counts k in the signal bin is given by

$$p = \sum_{n=k}^{\infty} \frac{\mu_\mathrm{back}^n e^{-\mu_\mathrm{back}}}{n!}\ , \tag{7.78}$$

where $k \in \mathbb{N}$. In order to find the sensitivity in axion parameter space, one has to evaluate the expectation value of p assuming a Poisson distribution for k with mean $\mu = \mu_\mathrm{back} + \mu_\mathrm{signal}$ for each possible value of the two relevant couplings. The

---

[60] This is only approximately true for the configurations $\mathrm{IAXO_r}$ and $\mathrm{IAXO_r^+}$ in table 7.5 for which an energy resolution equal to the FWHM of the emission peak was assumed.

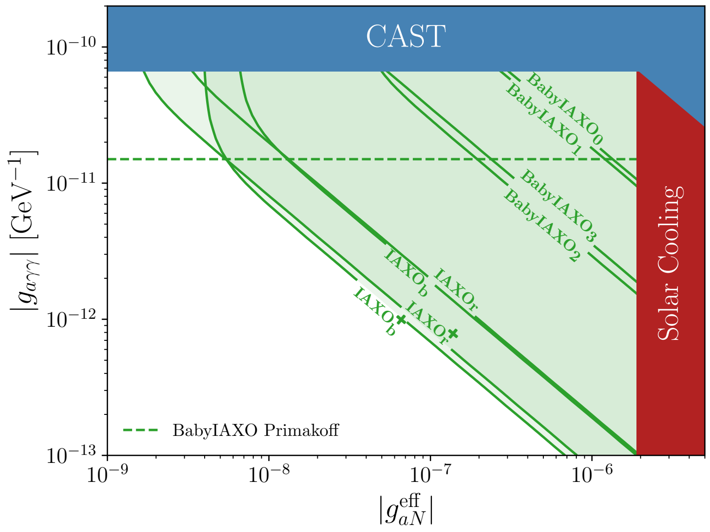

Figure 7.13. Model independent prediction for the sensitivity to the axion couplings for light axions ($\lesssim 20\,\text{meV}$). The different regions refer to the setups presented in table 7.5. The dark red region is the solar bound discussed in the text (see equation (7.70)). The dark blue region represents the latest CAST exclusion regions from searches for the Primakoff flux [208] and the ^{57}Fe peak [203]. The dashed horizontal green line indicates the expected sensitivity to the pure Primakoff flux.

experiment can be regarded as sensitive to a set of axion couplings if the expected p-value, $\langle p \rangle_k$, is smaller than 0.05.[61] The resulting sensitivity curves are plotted in figure 7.13.

Since the axion was assumed to be light or massless for the sensitivity estimates in figure 7.13, the parameter space of DFSZ and M1 models is not shown. This is because these models require large masses ($m_a \gtrsim 100\,\text{meV}$) at the couplings in question and helioscopes quickly lose sensitivity above $\sim 20\,\text{meV}$ (see figure 7.14). This problem could in principle be eased by filling the helioscope bore with a buffer gas [198].

The results in figure 7.13 show the potential of the various configurations to study areas of parameter space well beyond the solar bound and the region probed by CAST. The different green shaded areas indicate the expected sensitivity of the setups summarised in table 7.5. BabyIAXO would be most efficient when optimised optics are employed (setups labelled BabyIAXO$_{2,3}$). In this case, the total background is reduced by ~ 4 orders of magnitude with respect to the no-optics solution due to a small focal spot area a. BabyIAXO would also be able to extend its detection potential to regions of the parameter space below the Primakoff sensitivity if $g_{aN}^{\text{eff}} \gtrsim 10^{-7}$. This means it is possible that BabyIAXO could discover axions through the ^{57}Fe channel, before the Primakoff flux is detected. If, on the other hand, axions have couplings in the green shaded area above the Primakoff sensitivity line (dashed green), one might have the opportunity to extract both couplings

---

[61] A measurement with $p = 0.05$ would strictly speaking only amount to a 2σ anomaly. Nevertheless, it is common to define the sensitivity in this way because it coincides with the expected exclusion limits in the case of a null result.

and derive additional information about the underlying axion model.

The more optimistic IAXO and IAXO+ configurations can explore an even larger area of parameter space. A noteworthy feature of these exclusion curves is their behaviour at different values of $g_{a\gamma\gamma}$. At $g_{a\gamma\gamma} \lesssim 10^{-11}\,\mathrm{GeV}^{-1}$ the detector background dominates over the Primakoff background. Therefore the figure of merit in equation (7.73) becomes $f \propto \frac{\epsilon_o \epsilon_d \, g_{a\gamma\gamma}^2}{\sqrt{\Delta E_d b a}}$. With the parameters given in table 7.5, the configurations with minimised background slightly outperform the ones with optimised energy resolution in this area of parameter space (see figure 7.13). However, at $g_{a\gamma\gamma} \gtrsim 10^{-11}\,\mathrm{GeV}^{-1}$ the Primakoff background starts to play a role and eventually dominates. In this regime, the figure of merit is given by $f \propto \sqrt{\frac{\epsilon_o \epsilon_d}{\Delta E_d \kappa}}$. The detector background becomes negligible and the configurations with optimised energy resolution are significantly more sensitive than the ones with minimised background. Therefore, the ideal detector for the ^{57}Fe line crucially depends on the value of $g_{a\gamma\gamma}$. If BabyIAXO detects Primakoff axions, a detector with good energy resolution may be required to suppress the Primakoff background to the ^{57}Fe line. If, on the other hand, BabyIAXO only puts a stronger bound on $g_{a\gamma\gamma}$, the energy resolution becomes less important and the low-background detectors may be advantageous.

Including a finite axion mass

A finite axion mass can cause decoherence between the photon and axion wave functions inside the magnet bores and lead to a signal suppression. This effect is discussed in section 4.2 and becomes relevant for masses above $\sim 20\,\mathrm{meV}$ even though the exact value depends on the length of the magnet. The solar axion flux only starts to be affected by a finite axion mass of the order of keV, where IAXO would have already lost its sensitivity. Finite mass effects are therefore only relevant on the detection side. The expression for the conversion probability $P_{a\to\gamma}$ of axions into photons in the helioscope is given in equation (4.3). The suppression due to decoherence can be mitigated at the cost of some absorption by feeding a buffer gas into the bores [198, 199] (see equation (4.1)), which is why the sensitivity study to effectively massless (i.e. $m_a \lesssim 20\,\mathrm{meV}$) axions above serves as a good benchmark. Nonetheless, one can also explicitly investigate the sensitivity to massive axions without a buffer gas. To this end, it is helpful to assume that the background from Primakoff axions is negligible and that instead the detector background dominates. In this case, the background is independent of any axion properties and the signal depends on the product of the two couplings, $g_{a\gamma\gamma} g_{aN}^{\mathrm{eff}}$ as well as the mass. The statistical analysis is equivalent to the one in the massless case. The resulting sensitivity curves for all of the setups in table 7.5 are plotted in figure 7.14. The regions shaded in yellow indicate the coupling relations for the DFSZ, M1 and $\mathcal{T}_2^{(u)}$ models (see section 7.5.1).

The sensitivity curves are very similar to typical helioscope exclusion plots in the coupling vs. mass plane with two noteworthy exceptions. Because the ^{57}Fe line is highly-energetic at 14.4 keV, the transferred momentum q is smaller than for axions of the same mass from other solar processes. As a result the decoherence effect only becomes relevant at slightly higher masses in comparison to, for instance, Primakoff

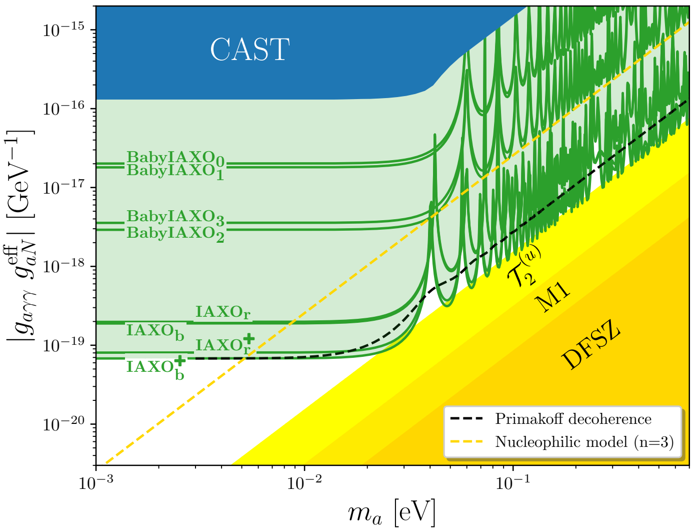

Figure 7.14. Model independent prediction of the IAXO sensitivity to the ^{57}Fe peak, assuming that axions are produced only through the axion coupling to nucleons. In contrast to figure 7.13, only the sensitivity of the various setups to the coupling combination $g_{a\gamma\gamma} g_{aN}^{\mathrm{eff}}$ is shown under the assumption that the Primakoff background is negligible. The oscillations at higher masses are due to the form factor in the conversion probability. For comparison, the effect of decoherence for a Primakoff spectrum is shown as a dashed black line. The dark blue region represents the rescaled CAST result [203]. The dark yellow region indicates the parameter space expected for the DFSZ model and brighter shades of yellow are used for flavour non-universal DFSZ models M1 and $\mathcal{T}_2^{(u)}$. The dashed yellow line shows the expected coupling for a nucleophilic QCD axion model of the kind presented in ref. [427], with $n = 3$. All models with $n > 3$ would be already accessible to BabyIAXO .

axions. To illustrate this effect, the expected sensitivity of the IAXO$^+_\mathrm{b}$ setup with a decoherence factor from a Primakoff spectrum is plotted as a dashed black line in figure 7.14. Furthermore, the oscillations of the form factor for large qL are clearly visible in the ^{57}Fe exclusion lines while they are washed out in the case of the broadband Primakoff spectrum.

In summary, the results in figures 7.13 and 7.14 show that already BabyIAXO will be able to study a large section of interesting axion parameter space, well beyond the region accessed by CAST, particularly in the configurations with optimised x-ray optics. The potential will be greatly improved with IAXO and IAXO+. Figure 7.14 also show representative QCD axion models like the DFSZ model and flavour non-universal M1 and $\mathcal{T}_2^{(u)}$ models discussed in section 7.5.1. The well-known KSVZ axion model does not appear because an accidental cancellation reduces its effective coupling to nucleons relevant in the ^{57}Fe transition. Although BabyIAXO is expected to have enough sensitivity to explore large sections of the parameter space for these models [123], a sizeable axion flux from ^{57}Fe transitions requires large axion masses, where BabyIAXO loses sensitivity. This problem can be eased with the use of a buffer gas [198], a technique already tested in CAST [208]. The sensitivities shown in the figures do not account for this option in BabyIAXO nor in its scaled up versions. A dedicated study may show if a buffer gas may allow to probe the M1 or other DFSZ-like models through the ^{57}Fe line in the near future.

Less minimal models for QCD axions may present larger couplings to nucleons and be better accessible through the ^{57}Fe channel. For example, the nucleophilic QCD axion models presented in ref. [427] (and shown in figure 7.14), have exponentially large couplings to nucleons[62] and are efficiently produced in ^{57}Fe transitions even at lower axion mass. One should nevertheless keep in mind that most of the region shown in the figures is in tension with astrophysical considerations, in particular, the bound from SN1987A as given in equation (7.72). Due to the uncertainties of these bounds, an independent confirmation would be highly desirable. That being said, the IAXO+ setup approaches a level of sensitivity comparable to supernova bound, showing a potential pathway for pushing beyond the astrophysical limits.

In conclusion, helioscopes of the next generation may offer a unique chance to probe an interesting range of the g_{aN}^{eff}-$g_{a\gamma\gamma}$ parameter space. BabyIAXO would benefit greatly from detector and optics designs optimised for the 14.4 keV line. In such a configuration, BabyIAXO has the potential to detect axions through the ^{57}Fe channel, which are too weakly coupled to photons to give a sizeable Primakoff flux (region below the dashed green line in figure 7.13). In a more likely scenario, a detection of axions through ^{57}Fe will be accompanied by a large signal from Primakoff axions, allowing to extract important information about its couplings to both photons and nucleons and thereby rule out a large number of possible UV completions.

7.6 Summary of results

The solar axion flux is one of the most expressive signatures of a QCD axion model. Several axion couplings enter the evaluation of the axion spectrum and in addition the axion mass can be inferred from the spectral shape of the conversion probability inside a helioscope. In light of this fact and the recently growing interest in solar axions, this chapter has revisited the processes generating the solar axion flux and assessed their uncertainties. This is motivated by the crucial influence that they can have when measuring QCD axion or axion-like particle properties.

The solar axion flux depends on solar properties and therefore on solar models. It was shown that the uncertainties of the axion flux resulting from the statistical variations within solar models of a given type are relatively modest in size. This shifts the focus to a better determination of systematic errors. One example is the metallicity problem, which yields flux uncertainties at the level of 10 % or more across a large energy range. Furthermore, calculations of the solar axion flux use several approximations, some of which may need to be improved in order to reduce the systematic theory errors. An important example is the partial degeneracy of

---

[62] The model A of ref. [427] has

$$g_{a\gamma\gamma} g_{aN}^{\text{eff}} \sim \frac{2^{2n} \alpha_{\text{EM}}}{2\pi f_a} \frac{m_n}{f_a} \sim 2^{2n+2} \cdot 10^{-17} \left(\frac{m_a}{\text{eV}}\right)^2 \text{GeV}^{-1} , \tag{7.79}$$

where $n+1$ is the number of Higgs doublets in the model. The parameter n is constrained to $n \lesssim 50$ in order to avoid sub-Planckian Landau poles [115].

electrons in the solar core (see ref. [233]), which was included in section 7.1.2 into the calculation of the Primakoff spectrum.

The ABC flux is also affected by uncertainties originating from variations between different opacity codes. Achieving percent-level accuracy here will likely require going beyond the approximation where the axion emission is directly proportional to the photon opacity and accounting for the differences between the emission of axions and photons from each atomic transition. Moreover, a better treatment of screening effects as well as non-trivial electron wave-functions (Sommerfeld enhancement) using a systematic, first-principles calculation would be very useful.

Beyond uncertainties in the standard Primakoff and ABC fluxes, there are additional production mechanisms such as the plasmon conversion in the Sun's magnetic field. The latter could add a noticeable flux contribution in the energy range of interest if the magnetic fields are at the higher end of the estimated ranges.

The inference of underlying axion model parameters is also limited by the current knowledge of low-energy axion parameters, such as the mass and the model-independent contribution to the axion-photon coupling (see e.g. refs. [54, 69]). Addressing these effects will be important in order to accurately determine the properties of axion models with helioscopes.

The Primakoff flux calculation is less affected by uncertainties, which makes it suitable for a study on the determination of axion model parameters. As a concrete example for such a study, a few benchmark cases of a potential future axion detection were considered in section 7.4. The known uncertainties on the solar Primakoff flux are sufficiently small that IAXO may indeed be able to identify and discriminate between different KSVZ axion models. It may even be possible to address solar physics questions by distinguishing between low-Z and high-Z solar models if the axion to photon coupling is sufficiently close to the current bound.

Finally, the potential of IAXO to detect the 14.4 keV line from nuclear transitions of ^{57}Fe was assessed in detail. Already BabyIAXO – with a few modifications – will be able to probe parameter space so far unexplored. It is even possible that a first axion detection will happen in the ^{57}Fe channel. All of this requires optimised detector and optics setups, which are outlined in section 7.5.3. It is, however, important to keep in mind that most of the region shown in figures 7.13 and 7.14 is in tension with astrophysical considerations, in particular, SN1987A (see equation (7.72)). Because these types of bounds are still under discussion [343] and are subject to some significant uncertainties, even a mere confirmation of the bound in a controlled setup would be a significant result. In addition, the IAXO+ setup shown in figure 7.13 approaches a level of sensitivity comparable to equation (7.72), showing a pathway for pushing beyond the astrophysical limits.

In conclusion, the solar axion flux, if detected, would likely be the most powerful observation for linking UV-complete solar axion models and their low-energy phenomenology. Any observation of all or some of the flux components discussed in this chapter would immediately rule out a large number of axion models. Moreover, a good agreement of an observations with the predicted spectrum calculated above would be a strong hint that an axion – and possibly a QCD axion – is being observed.

8 Conclusion and Outlook

The objective of this thesis is to investigate known and establish new connections between QCD axion models and their low-energy phenomenology. This is because axions stand out as one of the best motivated BSM hypotheses today since they can address several shortcomings of the SM at once and are an almost generic outcome of a large class of high-energy models. New experimental searches for axions are constantly being devised and the interplay between axion models and their low-energy phenomenology is crucial for investigating their expected sensitivity. These studies can furthermore help to optimise the experimental configurations for axions or axion-like particles.

Following the results in chapters 5 to 7, it becomes clear how the UV theory influences observations in the IR. The axion's periodicity is, for instance, a property that can only be understood by embedding the axion into a complete model. Crucially, the periodicity is protected at all energy scales and survives mixing with other pseudoscalars. One of the consequences of this property are quantisation conditions on specific combinations of axion couplings. Even though it seems as if loop corrections to low-energy effective couplings break this condition, it was shown how the periodicity is actually conserved at least when an appropriate operator basis is chosen and no periodicity violating approximations are applied. Because the periodicity is necessarily preserved, it can be an extremely helpful tool for analysing the validity of proposals involving an axion. This was demonstrated in two examples, the effective axion-photon coupling from mixings with pseudoscalar mesons and for the case of an axion-majoron model.

The periodicity chapter is a good example of the top-down approach in axion physics, while the second project presented in this thesis (chapter 6) proceeds in the opposite direction. It starts from the observation that the leading log approximation, often applied in EFT calculations, is not a good description for typical benchmark axion models. This immediately leads to the question whether this is true for all UV-complete axion models. By explicitly constructing a completion to an axion EFT without introducing new degrees of freedom below the PQ symmetry breaking scale, it was demonstrated that there are axion models with couplings of order unity to SM fermions and a cutoff at the PQ scale. The lessons from this discussion are widespread. First, it highlights the limitations of the EFT formalism, especially when there are other manifestations of a QCD axion model than the axion itself at a scale between experiment and PQ breaking. Second, a new model was devised. Due to its simplicity in a large energy range, it can be a useful benchmark for phenomenological studies. It was already shown that the two experiments NA62 and IAXO will be sensitive in an interesting range of parameter space. Finally, the work opens a discussion on which kinds of UV completions exist for a given set of EFT parameters. Even though a completion of the desired form was found in this case, it has a few possible issues like the quite complicated structure of the model, the potential appearance of additional flavour-changing couplings and non-perturbative Yukawas being required in some realisations to reproduce the top quark mass. A further investigation could reveal under which additional assumptions the set of axion EFTs which are the low-energy manifestations of reasonable UV models could be restricted.

Chapter 7 is focused on low-energy phenomena, namely all processes generating axions inside the Sun. In order to successfully relate an observation at low energies to an axion model, as many axion parameters as possible should be measured. To this end, it is beneficial to investigate the solar axion flux because it is unique in the amount of information encoded. Axion helioscopes like IAXO have potential access to the photon-, electron- and a combination of neutron- and proton-coupling as well as the mass of the axion – more than any other detection technique. The full potential of solar axion searches can only be exploited if the axion spectrum is known with the best-possible accuracy. Thus, this thesis combined a large number of previous works and several new ideas to provide *the* state-of-the-art calculation of the solar axion flux. For maximal flexibility, it is compatible with all publicly available standard solar models and opacity codes. The code which was developed for this purpose is publicly available and is already being used by the IAXO collaboration. The computationally efficient implementation also allowed to perform an in-depth analysis of associated uncertainties, using both MC simulations as well as estimations of systematic effects. A quantification of uncertainties is both important for statistical analysis of experimental data as well as for locating the parts of the calculation where improvements are most urgent. It was found that the Primakoff prediction already reaches percent-level accuracy. In contrast, the ABC flux is affected by much larger uncertainties related to both opacity codes and systematics of the calculation like the ones from partial degeneracy of electrons. Reducing this uncertainty in the future will require improvements of the calculation and in particular going beyond the direct relation between photon opacity and axion emission. As an example for how the uncertainties influence what can be learned from an axion observation, the simple case of a pure Primakoff flux was studied. With the updated calculation, it can already be possible to gain insights into solar properties from an axion detection. This shows the potential of axion physics and the possibility of a future axion precision age, in which the axion takes the role of a new messenger also revealing information about open questions unrelated to BSM physics.

An additional effort was made to study the sensitivity of IAXO to an isolated signal in the solar axion spectrum generated by a nuclear transition of the isotope ^{57}Fe. This signal could be very useful for ruling out a large number of possible axion models. Because the transition energy of 14.4 keV is larger than the typical ones of Primakoff or ABC axions, a detection of the ^{57}Fe peak requires modified experimental configurations. A number of options was discussed and it became clear that a combination of purposely built optics and high-resolution, low-background detectors will be required. Whether it is better to focus on optimising the energy resolution or on minimising the detector background depends on the strength of the axion-photon coupling and the corresponding Primakoff background to the ^{57}Fe peak. With such an optimised setup, IAXO could even make a first axion detection in the ^{57}Fe channel.

The summary above demonstrates how physics at different scales can be intricately linked to one another. A particularly illustrative example of these connections and how the different projects come together is discussed in section 7.4. By using the property of quantised couplings from chapter 5, one can draw additional information from an axion observations in the IR. In this case, it was used to break a near perfect degeneracy between the solar metallicity and the axion-photon coupling. A UV axion property may therefore be directly linked to the measurement of a poorly known solar parameter.

A large number of new axion experiments is planned to be built in the coming decades. This is going to be accompanied by new and ever more precise astrophysical observations, which can give new hints or bounds on axions. Collectively, these searches are going to probe the majority of well-motivated axion parameter space and finally give a definite answer to the question, whether the 45-year-old idea by Peccei and Quinn is realised in nature in one of its simplest forms. At the same time, future work in the direction of this thesis could be used to give predictions for high-energy physics from low-energy experiments. Suppose for instance that an axion with fermion couplings was discovered in experiments like IAXO or NA62. This could be interpreted as a hint for an extended Higgs sector because this is the typical way to generate this interaction in QCD axion models. Next generation collider experiments could then perform direct searches for exactly these heavy scalars. The future of particles physics, even at high energy, is therefore not exclusively relying on ever larger accelerators but progress could also be driven more and more by studying phenomena beyond the low-energy precision frontier.

Appendix

A CP violation in axion-fermion interactions

This appendix is based on results and arguments which were published in appendix A of ref. [2] before submission of this thesis. Material from this article is used in the following without further reference.

The most general description of axion-fermion interactions is given by

$$\frac{\partial_\mu a}{f_a}\sum_\chi \bar{\chi} C_\chi \gamma^\mu \chi\,. \tag{A.1}$$

Just like in section 3.4, the sum runs over all chiral multiplets $\chi = Q_L, L_L, u_R, d_R, e_R$. Under a CP transformation each multiplet transforms like

$$\chi \xrightarrow{\text{CP}} C\chi^* \quad \text{and} \quad \bar{\chi} \xrightarrow{\text{CP}} \bar{\chi}^* C \tag{A.2}$$

where C is the charge conjugation matrix, which is anti-symmetric, real, unitary and satisfies $(C\gamma_\mu)^T = C\gamma_\mu$. This is not to be confused with the coupling matrix C_χ, which is defined in generation space. Inserting this into the axion interaction term and reintroducing the generation indices, one can evaluate how the interaction term transforms under CP.

$$\begin{aligned}
\frac{\partial_\mu a}{f_a} C_{\chi,ij}\, \bar{\chi}_i\, \gamma^\mu\, \chi_j \quad \xrightarrow{\text{CP}} \quad & P_\mu^\nu \frac{\partial_\nu a}{f_a} C_{\chi,ij}\, \bar{\chi}_i^*\, C\gamma^\mu C\, \chi_j^* && \text{(A.3)}\\
=& P_\mu^\nu \frac{\partial_\nu a}{f_a} C_{\chi,ij}\, \chi_j^\dagger\, C\gamma^{\mu T} C\gamma^0\, \chi_i && \text{(A.4)}\\
=& P_\mu^\nu \frac{\partial_\nu a}{f_a} C_{\chi,ij}\, \bar{\chi}_j\, \gamma^0\gamma^\mu\gamma^0\, \chi_i && \text{(A.5)}\\
=& P_\mu^\nu \frac{\partial_\nu a}{f_a} C_{\chi,ij}\, P_\rho^\mu \bar{\chi}_j\, \gamma^\rho\, \chi_i && \text{(A.6)}\\
=& \frac{\partial_\mu a}{f_a} C_{\chi,ij}\, \bar{\chi}_j\, \gamma^\mu\, \chi_i && \text{(A.7)}
\end{aligned}$$

A basis for the gamma matrices was used in which $\gamma^{0\dagger} = \gamma^0$, $\gamma^{i\dagger} = -\gamma^i$ and $\gamma^0\gamma^\mu\gamma^0 = \gamma^{\mu\dagger}$. P denotes the parity matrix, $P = \text{diag}(1, -1, -1, -1)$. The calculation shows that CP maps C_χ onto its transpose. Since C_χ has to be hermitian due to hermiticity of the Lagrangian, the real part of C_χ is always symmetric and thereby CP conserving. The imaginary part is however anti-symmetric. Axion-fermion interactions are therefore CP violating if and only if the coupling matrix in generation space has complex off-diagonal entries. The same is true for interactions in a mass-diagonal basis where the same derivation applies. The vector and axial-vector couplings are just linear combinations of the chiral ones, so again CP violation occurs only in the case of complex off-diagonal entries.

It is noteworthy that the mass diagonalization procedure can in principle induce some degree of CP violation in the axion-quark interactions even though the matrices C_χ may have been real in the UV. This can be most directly seen in the relation between up- and down-type left-handed quark interactions, $c_{d,L} = V^\dagger\, c_{u,L}\, V$. Since the unitary CKM matrix V has been measured to have a non-vanishing CP violating phase, it is possible to get CP violating interactions in the case of flavour non-universal axion couplings. This CP violation is however fully determined by the CKM phase and does not originate from the axion sector.

B Variations of the new axion model

This appendix is based on results and arguments which were published in appendix B of ref. [2] before submission of this thesis. Material and figures from this article are used in the following without further reference.

The UV completion in equation (6.20) is not unique. Instead it is also possible to couple either or both of the up- and down-type fields to Φ^* instead of Φ. In total, this makes four possible ways to choose between Φ and Φ^* and the most general Lagrangian is given by

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

Because Φ transforms trivially under all of the SM gauge groups, this choice does not influence the gauge charges of any other fields. The difference appears in the relations of PQ charges of the different fermions. The terms in equation (B.1) enforce the conditions

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

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

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

The upper/lower sign corresponds to inserting Φ/Φ^* into the Lagrangian (B.1). As in the main text, the sum of the four equations above appears in the anomaly coefficient, because all SM and F quarks are in the fundamental representation of SU(3).

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

Hence, $|N| = 6$ as long as either Φ or Φ^* is chosen for both up- and down-type quarks. Otherwise, the two contributions cancel, there is no QCD anomaly and the pseudoscalar a is not a QCD axion. For this reason, this possibility is not pursued any further and PQ charges of up- and down-type fields are taken to be identical. The corresponding labels are dropped in the following.

After this, the only other choice is the PQ charge assignment of F quarks. These determine the possible origin of 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 chapter 6.

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

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

   This option is 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,

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

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

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

   The existence of these couplings slightly modifies the diagonalization of the mass matrix. Even though the SM quarks 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 $\epsilon^2 = \frac{v^2}{\langle\phi\rangle^2}$. This model will therefore not induce large flavour-changing effects in the IR.

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

$$\chi_{F_L} - \chi_{F_R} = 1 \quad \Rightarrow \quad \chi_{Q_L} - \chi_{q_R} = -2 \,. \tag{B.12}$$

   Both SM and F quarks have axial PQ charges and therefore tree-level axion couplings are present 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 (6.13) 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 it is done in the main text. In models 3 and 4, a full diagonalization including all SM and F quarks along the lines of what is done in appendix C is always necessary.

C Mass diagonalization and axion interactions

This appendix is based on results and arguments which were published in appendix C of ref. [2] before submission of this thesis. Material from this article is used in the following without further reference.

In the following, the full diagonalization procedure of chapter 6 is outlined. The resulting couplings of the axion and light and heavy quarks are derived. This includes some repetitions of intermediate steps and results which are also given in chapter 6.

Absorbing the axion field into the quarks

The mass diagonalization starts from the mass matrix in the UV-complete theory after PQ and electroweak symmetry breaking,

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

with the two expansion parameters

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

The transformation of quark fields is split into two parts, of which only the first one depends on the axion field. This is done by absorbing the axion dependence of the quark mass matrix into the right-handed quarks,

$$u_R \to e^{-\frac{ia}{\langle\phi\rangle}} u_R \, , \quad d_R \to e^{-\frac{ia}{\langle\phi\rangle}} d_R \, . \tag{C.3}$$

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

$$\bar{q}_R i\partial\!\!\!/ q_R \to \bar{q}_R i\partial\!\!\!/ q_R + \frac{\partial_\mu a}{\langle\phi\rangle} \bar{q}_R \gamma^\mu q_R \, , \tag{C.4}$$

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 [14, 49, 50, 58, 107]

$$\begin{aligned} \mathcal{L} \supset & - N \cdot \frac{\alpha_s}{8\pi \, \langle\phi\rangle} \, a \, G^a_{\mu\nu} \widetilde{G}^{a\mu\nu} - E \cdot \frac{\alpha_{\text{EM}}}{8\pi \, \langle\phi\rangle} \, a \, F_{\mu\nu} \widetilde{F}^{\mu\nu} \\ & + E \cdot \frac{\alpha_{\text{EM}}}{4\pi} \frac{s_W}{c_W} \frac{a}{\langle\phi\rangle} F_{\mu\nu} \widetilde{Z}^{\mu\nu} - E \cdot \frac{\alpha_{\text{EM}}}{8\pi} \frac{s_W^2}{c_W^2} \frac{a}{\langle\phi\rangle} Z_{\mu\nu} \widetilde{Z}^{\mu\nu} \, , \end{aligned} \tag{C.5}$$

In equation (C.4), the axion-dependent rotations were only applied 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

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

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

$T(R_f)$ is the Dynkin index of the SU(3) representation[63] and Q_f are the electric charges. In the new model, $N = 6$ and $E = 10$ or 16 depending on whether the leptons are also charged under PQ. The axion gluon coupling is normalized in the conventional manner,

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

with $N_{\text{DW}} \equiv N = 6$. With these definition, the gauge-boson coupling terms can be written as

$$\begin{aligned} \mathcal{L} \supset &- \frac{\alpha_s}{8\pi f_a} \, a \, G^a_{\mu\nu} \widetilde{G}^{a\mu\nu} - \frac{E}{N} \cdot \frac{\alpha_{\text{EM}}}{8\pi f_a} \, a \, F_{\mu\nu} \widetilde{F}^{\mu\nu} \\ &+ \frac{E}{N} \cdot \frac{\alpha}{4\pi} \frac{s_W}{c_W} \frac{a}{f_a} F_{\mu\nu} \widetilde{Z}^{\mu\nu} - \frac{E}{N} \cdot \frac{\alpha_{\text{EM}}}{8\pi} \frac{s_W^2}{c_W^2} \frac{a}{f_a} Z_{\mu\nu} \widetilde{Z}^{\mu\nu} \, . \end{aligned} \tag{C.9}$$

By comparing this to the EFT Lagrangian (3.21), one can identify $c_{gg} = 1$, $c_{WW} = 0$, $c_{\gamma\gamma} = E/N$, $c_{ZZ} = \tan(\theta_W)^2 \ E/N$ and $c_{\gamma Z} = -2\tan(\theta_W) \ E/N$.

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

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

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

Mass diagonalization

In order to fully diagonalize the matrices M^u and M^d, unitary transformations of left- and right-handed fields are required. As in the main text, the labels u/d and generation indices are dropped, but both types of quarks and all three generations of SM and F quarks are considered at every point.

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

$$U^\dagger M M^\dagger U = \Lambda^2, \tag{C.11}$$

---

[63] In the conventions of this thesis, $T(R_f) = 1$ for the fundamental representation, which differs from the published version in ref. [2] by a factor of 2.

where Λ^2 is a diagonal matrix with only real positive eigenvalues. Another unitary matrix is defined by $S = M^\dagger U \Lambda^{-1}$. The two matrices can be used to transform the left- and right-handed quark fields, respectively,

$$\begin{pmatrix} q_L \\ F_L \end{pmatrix} \to U \begin{pmatrix} q_L \\ F_L \end{pmatrix} \tag{C.12}$$

$$\begin{pmatrix} q_R \\ F_R \end{pmatrix} \to S \begin{pmatrix} q_R \\ F_R \end{pmatrix} \ . \tag{C.13}$$

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

$$U^\dagger M S = U^\dagger M M^\dagger U \Lambda^{-1} = \Lambda \ . \tag{C.14}$$

U can be found perturbatively, only using its property that it diagonalizes

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

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

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

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,

$$M_q^2 = \mathrm{diag}(m_{q_1}^2, m_{q_2}^2, m_{q_3}^2) \tag{C.17}$$

$$= \left[U_\delta^\dagger (\delta - \mu \xi^{-1} \mu^\dagger) U_\delta \, (\epsilon \epsilon')^2 + \mathcal{O}(\epsilon^3) \right] \frac{\langle \Sigma \rangle^2}{2} , \tag{C.18}$$

$$M_F^2 = \mathrm{diag}(m_{F_1}^2, m_{F_2}^2, m_{F_3}^2) \tag{C.19}$$

$$= \left[U_\xi^\dagger (\xi + \frac{(\epsilon\epsilon')^2}{2} (\mu^\dagger \mu \xi^{-1} + \xi^{-1} \mu^\dagger \mu)) U_\xi + \mathcal{O}(\epsilon^3) \right] \frac{\langle \Sigma \rangle^2}{2} . \tag{C.20}$$

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 sizeable corrections to the effective description in (6.13), it needs to be ensured 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, the scale separation between the mass of the F quarks and the PQ scale needs to be sufficiently large. This condition can be written as

$$\min_i m_{F_i}^2 = \min \mathrm{eig}(\xi) \, \langle \Sigma \rangle^2 \gg \langle \phi \rangle^2 \tag{C.21}$$

$$\Rightarrow \quad 1 \gtrsim \min \mathrm{eig}(\xi) = \min \mathrm{eig}(\lambda \lambda^\dagger + \epsilon'^2 \beta \beta^\dagger) \gg \epsilon'^2 \tag{C.22}$$

$$\Rightarrow \quad 1 \gtrsim \min \mathrm{eig}(\lambda \lambda^\dagger) \gg \epsilon'^2 \ . \tag{C.23}$$

'min eig' denotes the smallest eigenvalue of a matrix. In the second and third lines, the size of the eigenvalues are constrained by perturbativity. Note 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. It is therefore helpful to define another expansion parameter ϵ'' as

$$\epsilon''^2 = \frac{\epsilon'^2}{\min \operatorname{eig}(\lambda\lambda^\dagger)} \simeq \frac{\langle\phi\rangle^2}{\min_i m_{F_i}^2} \ . \tag{C.24}$$

When $\epsilon'' \ll 1$, ξ^{-1} can be expanded as

$$\begin{aligned} \xi^{-1} &= (\lambda\lambda^\dagger + \epsilon'^2\beta\beta^\dagger)^{-1} && \text{(C.25)} \\ &= \lambda^{\dagger-1}\ (1 + \epsilon'^2\lambda^{-1}\beta\beta^\dagger\lambda^{\dagger-1})^{-1}\ \lambda^{-1} && \text{(C.26)} \\ &= \lambda^{\dagger-1}\ \sum_n (-\epsilon'^2\lambda^{-1}\beta\beta^\dagger\lambda^{\dagger-1})^n\ \lambda^{-1} && \text{(C.27)} \\ &= \lambda^{\dagger-1}\ (1 - \epsilon'^2\lambda^{-1}\beta\beta^\dagger\lambda^{\dagger-1} + \mathcal{O}(\epsilon''^4))\ \lambda^{-1}\ , && \text{(C.28)} \end{aligned}$$

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

$$\begin{aligned} M_q^2 = \operatorname{diag}(m_{q_1}^2, m_{q_2}^2, m_{q_3}^2) &\simeq U_\delta^\dagger \alpha\lambda^{-1}\beta\beta^\dagger\lambda^{\dagger-1}\alpha^\dagger U_\delta\ \ \epsilon^2\epsilon'^4\frac{\langle\Sigma\rangle^2}{2} && \text{(C.29)} \\ &\simeq \mathcal{A}\mathcal{B}\mathcal{A}^\dagger\ \epsilon^2\epsilon'^4\frac{\langle\Sigma\rangle^2}{2}\ . && \text{(C.30)} \end{aligned}$$

In the last step, the coupling matrices $\mathcal{A}$ and $\mathcal{B}$ were defined as

$$\mathcal{A} = U_\delta^\dagger \mu\xi^{-1}U_\xi = U_\delta^\dagger\alpha\lambda^{-1}U_\xi + \mathcal{O}(\epsilon''^2) \quad \text{and} \quad \mathcal{B} = U_\xi^\dagger\beta\beta^\dagger U_\xi\ . \tag{C.31}$$

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 equation (C.4) does not transform trivially,

$$\frac{\partial_\mu a}{\langle\phi\rangle}\begin{pmatrix}\bar{q}_R & \bar{F}_R\end{pmatrix}\gamma^\mu\begin{pmatrix}\mathbb{1} & 0\\ 0 & 0\end{pmatrix}\begin{pmatrix}q_R\\ F_R\end{pmatrix} \to \frac{\partial_\mu a}{\langle\phi\rangle}\begin{pmatrix}\bar{q}_R & \bar{F}_R\end{pmatrix}\gamma^\mu S^\dagger\begin{pmatrix}\mathbb{1} & 0\\ 0 & 0\end{pmatrix} S\begin{pmatrix}q_R\\ F_R\end{pmatrix}\ . \tag{C.32}$$

To leading order in ϵ, S is given by

$$\begin{aligned} S =& M^\dagger U\Lambda^{-1} && \text{(C.33)} \\ =& \frac{\langle\Sigma\rangle}{\sqrt{2}}\Bigg[\begin{pmatrix} \epsilon\epsilon'^2\,\beta^\dagger\xi^{-1}\mu^\dagger U_\delta & \epsilon'\,\beta^\dagger(1 - \frac{(\epsilon\epsilon')^2}{2}\xi^{-1}\mu^\dagger\mu\xi^{-1})U_\xi \\ \epsilon\epsilon'\,(\lambda^\dagger\xi^{-1}\mu^\dagger - \alpha^\dagger)U_\delta & (\lambda^\dagger + (\epsilon\epsilon')^2(\alpha^\dagger\mu\xi^{-1} - \frac{1}{2}\lambda^\dagger\xi^{-1}\mu^\dagger\mu\xi^{-1}))U_\xi \end{pmatrix} \\ & \qquad + \mathcal{O}(\epsilon^3)\Bigg]\Lambda^{-1}\ , && \text{(C.34)} \end{aligned}$$

from which the relevant coupling matrix can be determined to leading order in ϵ as

$$\begin{aligned}
&\frac{\not{\partial} a}{\langle\phi\rangle} S^\dagger \begin{pmatrix} \mathbb{1} & 0 \\ 0 & 0 \end{pmatrix} S \\
&= \frac{\not{\partial} a}{\langle\phi\rangle} \Lambda^{-1} \frac{\langle\Sigma\rangle^2}{2} \begin{pmatrix} \epsilon^2\epsilon'^4 U_\delta^\dagger \mu \xi^{-1} \beta\beta^\dagger \xi^{-1} \mu^\dagger U_\delta & \epsilon\epsilon'^3 U_\delta^\dagger \mu \xi^{-1} \beta\beta^\dagger U_\xi \\ \epsilon\epsilon'^3 U_\xi^\dagger \beta\beta^\dagger \xi^{-1} \mu^\dagger U_\delta & \epsilon'^2 U_\xi^\dagger \beta\beta^\dagger U_\xi \end{pmatrix} \Lambda^{-1} && \text{(C.35)} \\
&= \frac{\not{\partial} a}{\langle\phi\rangle} \Lambda^{-1} \frac{\langle\Sigma\rangle^2}{2} \begin{pmatrix} \epsilon^2\epsilon'^4 \mathcal{A}\mathcal{B}\mathcal{A}^\dagger & \epsilon\epsilon'^3 \mathcal{A}\mathcal{B} \\ \epsilon\epsilon'^3 \mathcal{B}\mathcal{A}^\dagger & \epsilon'^2 \mathcal{B} \end{pmatrix} \Lambda^{-1} && \text{(C.36)} \\
&\simeq (\not{\partial} a) \begin{pmatrix} 1/\langle\phi\rangle & v\,\epsilon'\, M_q^{-1} \mathcal{A}\,\mathcal{B} M_F^{-1}/2 \\ v\,\epsilon' M_F^{-1} \mathcal{B}\,\mathcal{A}^\dagger M_q^{-1}/2 & \langle\phi\rangle\, M_F^{-1} \mathcal{B} M_F^{-1}/2 \end{pmatrix}, && \text{(C.37)}
\end{aligned}$$

where only the leading-order term in ϵ'' was kept in the last step as in equation (C.29), 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, it was also used that $\Lambda^{-1} = \mathrm{diag}(M_q^{-1}, M_F^{-1})$. With this, the axion couplings can be written as

$$\begin{aligned}
\mathcal{L} \supset& \frac{\langle\phi\rangle}{2} \begin{pmatrix} \bar{q}_R & \bar{F}_R \end{pmatrix} (\not{\partial} a)\, \Lambda^{-1} \begin{pmatrix} (\epsilon\epsilon')^2 \mathcal{A}\mathcal{B}\mathcal{A}^\dagger & \epsilon\epsilon' \mathcal{A}\mathcal{B} \\ \epsilon\epsilon' \mathcal{B}\mathcal{A}^\dagger & \mathcal{B} \end{pmatrix} \Lambda^{-1} \begin{pmatrix} q_R \\ F_R \end{pmatrix} && \text{(C.38)} \\
\simeq& \frac{1}{\langle\phi\rangle} \bar{q}_R\, (\not{\partial} a)\, q_R + \frac{\langle\phi\rangle}{2} \bar{F}_R (M_F^{-1} \mathcal{B} M_F^{-1})\, (\not{\partial} a)\, F_R \\
&+ \left(\frac{v}{2}\, \epsilon' \bar{q}_R (M_q^{-1} \mathcal{A}\,\mathcal{B} M_F^{-1})\, (\not{\partial} a)\, F_R + \text{h.c.} \right), && \text{(C.39)}
\end{aligned}$$

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

Tree-level contributions to $s \to d + a$

It was shown in the previous section that, at leading order, the new model only includes flavour-diagonal tree-level couplings between the axion and SM quarks. In this case, the $s \to d + a$ process is only induced at one-loop level which is also the scenario discussed in the main text. However, one cannot exclude the possibility that higher orders in the expansion of both ϵ and ϵ'', which induce non-diagonal coupling structures, become relevant. These flavour-violating couplings would trigger the $s \to d + a$ decay already at tree level. It therefore becomes important to quantify them and to check if they restrict the range of values of the expansion parameters.

To start with, one can quantify 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 $\mathrm{BR}(K^+ \to \pi^+ + a) < 7.3 \times 10^{-11}$ [322]. Generally, a coupling of the form

$$\mathcal{H}_{s\to da} = (\partial_\mu a)\, \bar{d} h_{ds}^S \gamma^\mu (1 + \gamma^5) s + \text{h.c.} \tag{C.40}$$

results in the decay width of equation (6.4)

$$\Gamma(K^+ \to \pi^+ a) = \frac{|h_{ds}^S|^2}{16\pi m_{K^+}^3} (m_{K^+}^2 - m_{\pi^+}^2)^2 \lambda^{1/2}(m_{K^+}^2,\, m_{\pi^+}^2,\, m_a^2)\, f_+^2(m_a^2)\,, \tag{C.41}$$

which leads to

$$|h^S_{ds}| \lesssim 1.5 \times 10^{-12}\,\mathrm{GeV}^{-1}\,. \tag{C.42}$$

It is straightforward to conclude that the expansion in ϵ is safe: The parameter $\epsilon = v/\langle\phi\rangle$ is $\mathcal{O}(10^{-4})$ for $f_a \sim 10^6\,\mathrm{GeV}$, which is the region where NA62 is sensitive. The leading-order axial-vector coupling between quarks and the axion in equation (C.39) therefore corresponds to $(2\langle\phi\rangle)^{-1} \sim (12\cdot f_a)^{-1} \sim 10^{-7}\,\mathrm{GeV}^{-1}$. 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 the higher-order terms in S in equation (C.34) were included. Therefore, the NLO coupling is parametrically suppressed by $10^{-7}\cdot 10^{-8}\,\mathrm{GeV}^{-1} = 10^{-15}\,\mathrm{GeV}^{-1}$, which is sufficiently far away from the E787 bound quoted in equation (C.42) and is also out of reach for NA62.

Next, it is necessary to quantify how small ϵ'' has to be in order to avoid significant flavour-violating axion couplings at tree level. By expanding ξ^{-1} in ϵ'' as in equation (C.25), one can write S schematically as

$$S = \begin{pmatrix} \text{unitary} + \mathcal{O}(\epsilon''^2) & \mathcal{O}(\epsilon'') \\ \mathcal{O}(\epsilon'') & \text{unitary} + \mathcal{O}(\epsilon''^2) \end{pmatrix}\,, \tag{C.43}$$

and estimate the parametric size of the coupling structure in equation (C.32) as

$$S^\dagger \begin{pmatrix} \mathbb{1} & 0 \\ 0 & 0 \end{pmatrix} S = \begin{pmatrix} \mathbb{1} + \mathcal{O}(\epsilon''^2) & \mathcal{O}(\epsilon'') \\ \mathcal{O}(\epsilon'') & \mathcal{O}(\epsilon''^2) \end{pmatrix}\,. \tag{C.44}$$

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

$$|h^{S,\mathrm{treelevel}}_{ds}| = \frac{1}{2\langle\phi\rangle}\left[S^\dagger \begin{pmatrix} \mathbb{1} & 0 \\ 0 & 0 \end{pmatrix} S\right]_{ds} \lesssim \frac{\epsilon''^2}{2\langle\phi\rangle}. \tag{C.45}$$

To avoid the constraint in (C.42), 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.

At this point, it is important to note two more interesting facts about this new model. First, flavour-violating axion couplings at tree level are not necessarily a problem but could also be considered an interesting feature of the model. Only because it was initially set up to UV-complete the model in (6.13) and not to build an axion-flavour model, the main part of this thesis is confined 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 $\epsilon' = 0.2$, one 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 equation (C.36) at all orders in ϵ'', as the inversion of ξ in (C.25) 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'',

$$M_q^2 = \left[U_\delta^\dagger(\delta - \mu\xi^{-1}\mu^\dagger)U_\delta\,(\epsilon\epsilon')^2 + \mathcal{O}(\epsilon^3)\right]\frac{\langle\Sigma\rangle^2}{2} \tag{C.46}$$

$$= \left[U_\delta^\dagger(\alpha\alpha^\dagger - \frac{1}{1+\epsilon'^2}\alpha\alpha^\dagger)U_\delta\,(\epsilon\epsilon')^2 + \mathcal{O}(\epsilon^3)\right]\frac{\langle\Sigma\rangle^2}{2} \tag{C.47}$$

$$= \frac{1}{1+\epsilon'^2}\mathcal{A}\mathcal{B}\mathcal{A}^\dagger\epsilon^2\epsilon'^4\frac{\langle\Sigma\rangle^2}{2} + \mathcal{O}(\epsilon^3). \tag{C.48}$$

Also note that with theses choices of couplings, the corrections to the relation in equation (C.29), which are crucial to obtain the effective value of h_{ds}, are of the order of $\epsilon''^2 = \epsilon'^2 = 4\,\%$. This means that the results and discussions of the $K^+ \to \pi^+ + a$ decay are still valid. Nevertheless, α is already very close to the perturbativity limit of Yukawa couplings in this explicit realisation .

Electroweak interactions

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

$$U = \begin{pmatrix} A & B \\ C & D \end{pmatrix}, \tag{C.49}$$

where each block is a 3×3 matrix. Note, however, that unitarity of U does not imply unitarity of any of the blocks. Because the transformation above mixes different representations of $\mathrm{SU}(2)\times\mathrm{U}(1)_\mathrm{Y}/\mathrm{U}(1)_\mathrm{EM}$, it needs to be checked how the interactions with W and Z bosons are modified. Starting with the Z bosons,

$$\begin{aligned} &\begin{pmatrix} \bar{q} \\ \bar{F} \end{pmatrix}^T \gamma_\mu Z^\mu \frac{-g}{\cos(\theta_W)}\left[\pm\frac{1}{2}\begin{pmatrix} \mathbb{1} & 0 \\ 0 & 0 \end{pmatrix} P_L - Q\sin^2(\theta_W)\begin{pmatrix} \mathbb{1} & 0 \\ 0 & \mathbb{1} \end{pmatrix}\right]\begin{pmatrix} q \\ F \end{pmatrix} \\ \to &\begin{pmatrix} \bar{q} \\ \bar{F} \end{pmatrix}^T \gamma_\mu Z^\mu \frac{-g}{\cos(\theta_W)}\left[\pm\frac{1}{2}\begin{pmatrix} A^\dagger A & A^\dagger B \\ B^\dagger A & B^\dagger B \end{pmatrix} P_L - Q\sin^2(\theta_W)\begin{pmatrix} \mathbb{1} & 0 \\ 0 & \mathbb{1} \end{pmatrix}\right]\begin{pmatrix} q \\ F \end{pmatrix}. \end{aligned} \tag{C.50}$$

In this expression, the upper (lower) sign refers to up- (down-) type quarks and Q is the electromagnetic charge. 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 the perturbative result of U in equation (C.16), one can find the Z-interactions at leading order in ϵ to be

$$\begin{aligned} \mathcal{L} \supset \begin{pmatrix} \bar{q} \\ \bar{F} \end{pmatrix}^T \gamma_\mu Z^\mu \frac{-g}{\cos(\theta_W)}\Bigg[&\pm \frac{1}{2}\begin{pmatrix} \mathbb{1} & -\epsilon\epsilon'\mathcal{A} \\ -\epsilon\epsilon'\mathcal{A}^\dagger & (\epsilon\epsilon')^2\mathcal{A}^\dagger\mathcal{A} \end{pmatrix} P_L \\ &- Q\sin^2(\theta_W)\begin{pmatrix} \mathbb{1} & 0 \\ 0 & \mathbb{1} \end{pmatrix}\Bigg]\begin{pmatrix} q \\ F \end{pmatrix}. \end{aligned} \tag{C.51}$$

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

The analogous computation has to be performed for the W interactions, but because these mix up- and down-type quarks, the corresponding labels are reintroduced. The W terms become

$$\frac{-g}{\sqrt{2}} \begin{pmatrix} \bar{u}_L \\ \bar{d}_L \\ \bar{F}^u_L \\ \bar{F}^d_L \end{pmatrix}^T \gamma^\mu \left[W^+_\mu \begin{pmatrix} 0 & \mathbb{1} & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} + W^-_\mu \begin{pmatrix} 0 & 0 & 0 & 0 \\ \mathbb{1} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \right] \begin{pmatrix} u_L \\ d_L \\ F^u_L \\ F^d_L \end{pmatrix}$$

$$\longrightarrow \frac{-g}{\sqrt{2}} \begin{pmatrix} \bar{u}_L \\ \bar{d}_L \\ \bar{F}^u_L \\ \bar{F}^d_L \end{pmatrix}^T \gamma^\mu \left[W^+_\mu \begin{pmatrix} 0 & A^{u\dagger}A^d & 0 & A^{u\dagger}B^d \\ 0 & 0 & 0 & 0 \\ 0 & B^{u\dagger}A^d & 0 & B^{u\dagger}B^d \\ 0 & 0 & 0 & 0 \end{pmatrix} \right.$$

$$\left. + W^-_\mu \begin{pmatrix} 0 & 0 & 0 & 0 \\ A^{d\dagger}A^u & 0 & A^{d\dagger}B^u & 0 \\ 0 & 0 & 0 & 0 \\ B^{d\dagger}A^u & 0 & B^{d\dagger}B^u & 0 \end{pmatrix} \right] \begin{pmatrix} u_L \\ d_L \\ F^u_L \\ F^d_L \end{pmatrix} . \tag{C.52}$$

The SM CKM matrix V is identified as

$$V = A^{u\dagger}A^d = U^{u\dagger}_\delta U^d_\delta + \mathcal{O}(\epsilon^2) \ , \tag{C.53}$$

which unlike in the SM does not have to be unitary, but non-unitarity only appears at order ϵ^2. The results for A and B can be inserted to get

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

$$\left. + \ W^-_\mu \begin{pmatrix} 0 & 0 & 0 & 0 \\ V^\dagger & 0 & -\epsilon\epsilon' V^\dagger \mathcal{A}_u & 0 \\ 0 & 0 & 0 & 0 \\ -\epsilon\epsilon' \mathcal{A}_d^\dagger V^\dagger & 0 & (\epsilon\epsilon')^2 \mathcal{A}_d^\dagger V^\dagger \mathcal{A}_u & 0 \end{pmatrix} \right] \begin{pmatrix} u_L \\ d_L \\ F^u_L \\ F^d_L \end{pmatrix} , \tag{C.54}$$

which is the expression used in the main text for the calculation of the $s \to d + a$ amplitude in the new model.

Radial modes of H and Φ

Up to this point, the radial modes of the Higgs field H and of the PQ field Φ were neglected, but they are of course present in the UV model and can in principle be involved in observable processes. In order to find their coupling structure, it is easiest to think of them as additional non-constant terms in the mass matrix,

$$\mathcal{L} \supset - \begin{pmatrix} \bar{q}_L & \bar{F}_L \end{pmatrix} M_{\text{rad}} \begin{pmatrix} q_R \\ F_R \end{pmatrix} + \text{h.c.}, \tag{C.55}$$

where

$$M_{\text{rad}} = \begin{pmatrix} 0 & \alpha \frac{h}{\sqrt{2}} \\ \beta \frac{\phi}{\sqrt{2}} & 0 \end{pmatrix} \,. \tag{C.56}$$

The unitary transformations then result in

$$M_{\text{rad}} \to U^\dagger M_{\text{rad}} S \,. \tag{C.57}$$

The full expressions, where U and S are expressed in terms of the UV parameters, are quite lengthy. That is why they are given separately for each 3×3 block below.

$$\begin{aligned}
[U^\dagger M_{\text{rad}} S]_{qq} =& \frac{\langle\Sigma\rangle}{2} \Big(\Big(\epsilon^2 \epsilon'^3 \mathcal{A}\mathcal{B}\mathcal{A}^\dagger \phi + \mathcal{O}(\epsilon^4) \Big) \\
&\qquad - \Big(\epsilon\epsilon' U_\delta^\dagger \alpha (\lambda^\dagger \xi^{-1} \mu^\dagger - \alpha^\dagger) U_\delta \, h + \mathcal{O}(\epsilon^3) \Big) \Big) M_q^{-1} && \text{(C.58)} \\
=& \frac{\langle\Sigma\rangle}{2} \Big(\epsilon^2 \epsilon'^3 \mathcal{A}\mathcal{B}\mathcal{A}^\dagger \phi + \frac{2}{\epsilon\epsilon' \langle\Sigma\rangle^2} M_q^2 \, h \Big) M_q^{-1} && \text{(C.59)} \\
\simeq& \frac{M_q}{\langle\phi\rangle} \phi + \frac{M_q}{v} h \,, && \text{(C.60)}
\end{aligned}$$

$$[U^\dagger M_{\text{rad}} S]_{qF} = \frac{\langle\Sigma\rangle}{2} \Big(\Big(\epsilon\epsilon'^2 \mathcal{A}\mathcal{B}\phi + \mathcal{O}(\epsilon^3) \Big) - \Big(\mathcal{C} \, h + \mathcal{O}(\epsilon^2) \Big) \Big) M_F^{-1} \,, \tag{C.61}$$

$$\begin{aligned}
[U^\dagger M_{\text{rad}} S]_{Fq} =& \frac{\langle\Sigma\rangle}{2} \Big(\Big(\epsilon\epsilon'^2 \mathcal{B}\mathcal{A}^\dagger \phi + \mathcal{O}(\epsilon^3) \Big) \\
&\qquad + \Big(\epsilon^2 \epsilon'^2 \mathcal{A}^\dagger U_\delta^\dagger \alpha (\lambda^\dagger \xi^{-1} \mu^\dagger - \alpha^\dagger) U_\delta h + \mathcal{O}(\epsilon^3) \Big) \Big) M_q^{-1} && \text{(C.62)} \\
=& \frac{\langle\Sigma\rangle}{2} \Big(\epsilon\epsilon'^2 \mathcal{B}\mathcal{A}^\dagger \phi - \frac{2}{\langle\Sigma\rangle^2} \mathcal{A}^\dagger M_q^2 \, h \Big) M_q^{-1} \,, && \text{(C.63)}
\end{aligned}$$

$$[U^\dagger M_{\text{rad}} S]_{FF} = \frac{\langle\Sigma\rangle}{2} \Big(\Big(\epsilon' \mathcal{B} \, \phi + \mathcal{O}(\epsilon^2) \Big) + \Big(\epsilon\epsilon' \mathcal{A}^\dagger \mathcal{C} \, h + \mathcal{O}(\epsilon^3) \Big) \Big) M_F^{-1} \,, \tag{C.64}$$

where $\mathcal{C}$ was defined as $\mathcal{C} = U_\delta^\dagger \alpha \lambda^\dagger U_\xi \simeq 2 \, \mathcal{A} \, M_F^2 / \langle\Sigma\rangle^2$. First, note that modifications of the Higgs coupling to SM quarks always appear together with additional factors of ϵ^2 and are thus negligible. Second, any additional contributions to $s \to d + a$ involving ϕ must involve internal F quarks and are therefore proportional to ϵ''^2. Consequently, any of these terms becomes negligible in the $\epsilon'' \ll 1$ limit (as long as this suppression also compensates for possible enhancements originating from the hierarchical coupling matrices). This is particularly important for the otherwise unsuppressed coupling in equation (C.63). But even without a significant ϵ'' suppression, the flavour violation can be small for some realisations of the coupling matrices α and β. For instance, in the realisation proposed above equation (C.46), the ϕ- and Higgs-induced loop processes become flavour-diagonal because the amplitude has the same flavour structure as the mass matrix.

D Counterterm contribution to the kaon decay rate

This appendix is based on results and arguments which were published in appendix D of ref. [2] before submission of this thesis. Material and figures from this article are used in the following without further reference.

It was acknowledged in the main text that counterterm contributions are relevant for the computation of the $s \to d + a$ transition. For instance, the diagram on the right-hand side of figure 6.4 is a counterterm contribution from the renormalization of quark fields. The important contributions all stem from the renormalization of the lighter SM quarks because the one related to heavy F fields are of higher order in ϵ''. Also SM quarks encounter divergent loop integrals involving the new heavy fields, but these are again parametrically suppressed by $\mathcal{O}(\epsilon''^2)$. Cancelling all divergences, including these higher-order terms, would require a more complete renormalization discussion compared to the one below, where only the leading-order processes in ϵ''^2 are considered.

The bare quark fields are labelled b and can be related to the renormalized fields as in ref. [428]

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

For the $s \to d + a$ transition only the down-type fields are relevant. The renormalization constants are determined by demanding that the one-loop $W^\pm$ contribution to $s \to d$ and the counterterm cancel each other [329], as depicted in figure D.1. The renormalization constants are therefore of quadratic order in the weak coupling constant.

The renormalized quark fields from equation (D.1) have to be inserted into the right-handed derivative axion-fermion interactions. This leads to the terms

$$\mathcal{L} = \frac{1}{\langle\phi\rangle}\bar{q}_R \left(\partial\!\!\!/ a\right) q_R + \frac{1}{2\langle\phi\rangle}\bar{d}_R \delta Z_{sd}^{R*} \left(\partial\!\!\!/ a\right) s_R + \frac{1}{2\langle\phi\rangle}\bar{d}_R \left(\partial\!\!\!/ a\right) \delta Z_{ds}^{R} s_R + \text{h.c.} . \tag{D.2}$$

Because δZ is expected to have non-vanishing off-diagonal entries, the new terms induce $s \to d + a$ at tree level. These processes are of order $g^2/\langle\phi\rangle$, just like the one-loop FCNCs in equation (6.45). Hence, there is no reason to neglect the counterterm contribution at this stage. As is known in the literature, the explicit renormalization

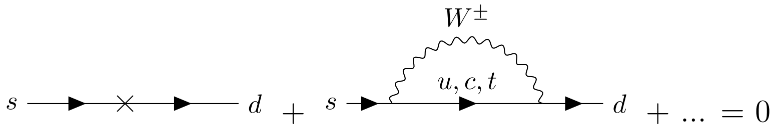

Figure D.1. On-shell renormalization condition of SM quarks. Only the relevant diagrams for determining the renormalization constants contributing to $s \to d$ are depicted. The dots stand for higher-order diagrams in the UV model, such as Z boson-induced flavour changes or diagrams with heavy F quarks.

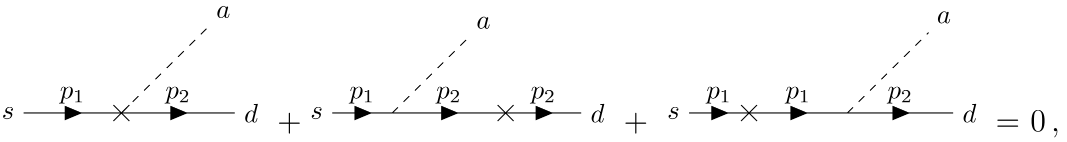

Figure D.2. Three different kinds of counterterm induced diagrams for $s \to d + a$ exactly add up to zero. The relation is proved below.

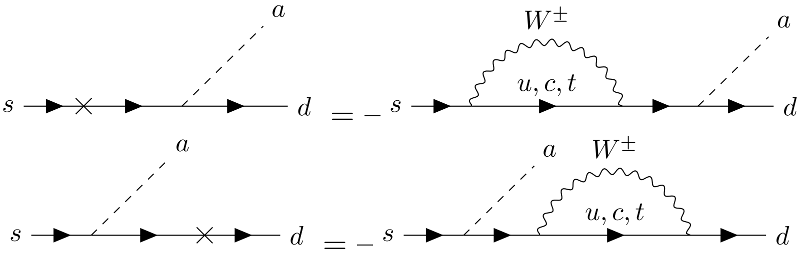

Figure D.3. Counterterm insertions are related to the $W^\pm$ loops for $s \to d + a$ because of the renormalization condition in figure D.1.

calculation does not need to be performed. Instead, the diagrammatic equation in figure D.2 can be shown to hold without knowing the precise expressions for the renormalization constants. In addition, one can also conclude the two relations given in figure D.3 from the renormalization condition shown in figure D.1.

By combining the three identities in figures D.1 to D.3, one obtains the key relation depicted in figure D.4. It allows to replace the counterterm contributions with additional self-energy diagrams, where the ALP is emitted from the external down-type quark legs [429], i.e. the diagrams on the right-hand side. It is therefore most straightforward to compute the additional loop diagrams instead of taking a detour to first compute the renormalization constants. When performing the calculations, one notices that the diagrams on the right-hand side of figure D.4 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. This is the result used in section 6.1.

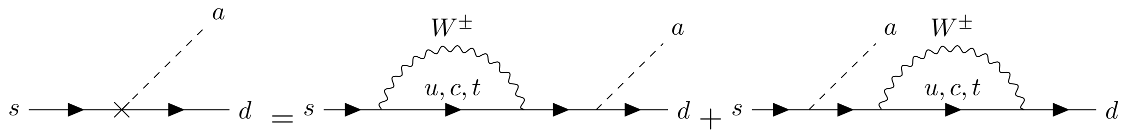

Figure D.4. Key relation between $s \to d+a$ counterterm contribution and ALP emission from external legs. This allows to skip the explicit evaluation of counterterms and instead add two more loop diagrams to the calculation of the matrix element.

Proof of the relation in figure D.2

Derivative axion couplings to right-handed quarks can be written in their most general form as

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

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. $G_Q = \mathrm{diag}(g_d,\, g_s,\, g_b)$. Only interactions with right-handed fermionic fields are considered here, but the derivation works equivalently for left-handed fields. Inserting the renormalized quark fields from equation (D.1), one obtains

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

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

The kinetic and mass terms are modified as

$$\mathcal{L} \supset i\bar{q}^b \partial\!\!\!/ q^b - \bar{q}^b M q^b \tag{D.6}$$

$$= i\bar{q}_R^b \partial\!\!\!/ q_R^b + i\bar{q}_L^b \partial\!\!\!/ q_L^b - \bar{q}_R^b M q_L^b - \bar{q}_L^b M q_R^b \tag{D.7}$$

$$= i\bar{q}_R \partial\!\!\!/ q_R + i\bar{q}_L \partial\!\!\!/ q_L - \bar{q}_R M q_L - \bar{q}_L M q_R + \mathcal{L}_C \,, \tag{D.8}$$

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

$$\begin{aligned}\mathcal{L}_C =& \frac{i}{2}\bar{q}_R (\delta Z^R)^\dagger \partial\!\!\!/ q_R + \frac{i}{2}\bar{q}_R \partial\!\!\!/ \, \delta Z^R q_R + \frac{i}{2}\bar{q}_L (\delta Z^L)^\dagger \partial\!\!\!/ q_L + \frac{i}{2}\bar{q}_L \partial\!\!\!/ \, \delta Z^L q_L \\ &- \frac{1}{2}\bar{q}_R (\delta Z^R)^\dagger M q_L - \frac{1}{2}\bar{q}_R M \, \delta Z^L q_L - \frac{1}{2}\bar{q}_L (\delta Z^L)^\dagger M q_R - \frac{1}{2}\bar{q}_L M \delta Z^R q_R \,.\end{aligned} \tag{D.9}$$

The left diagram in figure D.2 can be computed by just using equation (D.4),

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

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

The middle diagram involves an internal propagator and can be evaluated applying equations (D.4) and (D.9) to find

$$\begin{aligned}
i\mathcal{M}_2 =&\bar{u}_d(p_2)\frac{i}{2}\bigg\{ \Big[\Big(\delta Z^{R*}_{sd} + \delta Z^{R}_{ds}\Big)\not{p}_2 - \Big(m_d\delta Z^{R}_{ds} + m_s\delta Z^{L*}_{sd}\Big)\Big] P_R \\
&\qquad + \Big[\Big(\delta Z^{L*}_{sd} + \delta Z^{L}_{ds}\Big)\not{p}_2 - \Big(m_d\delta Z^{L}_{ds} + m_s\delta Z^{R*}_{sd}\Big)\Big] P_L\bigg\} \\
&\times \frac{i(\not{p}_2 + m_s)}{m_d^2 - m_s^2}\, i(p_1 - p_2)_\mu i g_s \gamma^\mu P_R u_s(p_1)
\end{aligned} \tag{D.11}$$

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

where the equation of motion $\bar{u}_d(p_2)\not{p}_2 = \bar{u}_d(p_2) m_d$ was applied. Using it once more for $\not{p}_2$ in the propagator yields

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

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

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

This precisely cancels the g_s term in equation (D.10). The terms involving δZ^L cancel due to $P_R P_L = 0$.

The right-hand diagram in figure D.2 is computed in the identical manner. The calculation reads

$$
\begin{aligned}
i\mathcal{M}_3 =& \bar{u}_d(p_2) i(p_1 - p_2)_\mu \, i g_d \, \gamma^\mu P_R \frac{i(\not{p}_1 + m_d)}{m_s^2 - m_d^2} \frac{i}{2} \\
& \Big\{ \Big[\Big(\delta Z_{sd}^{R*} + \delta Z_{ds}^{R} \Big) \not{p}_1 - \Big(m_d \delta Z_{ds}^{R} + m_s \delta Z_{sd}^{L*} \Big) \Big] \, P_R \\
& \quad + \Big[\Big(\delta Z_{sd}^{L*} + \delta Z_{ds}^{L} \Big) \not{p}_1 - \Big(m_d \delta Z_{ds}^{L} + m_s \delta Z_{sd}^{R*} \Big) \Big] \, P_L \Big\} u_s(p_1) \qquad \text{(D.16)} \\
=& \frac{-i(p_1 - p_2)_\mu}{2} \, i g_d \, \bar{u}_d(p_2) \gamma^\mu P_R \frac{\not{p}_1 + m_d}{m_s^2 - m_d^2} \\
& \quad \times \Big\{ \delta Z_{ds}^{R} m_s P_L - m_d \delta Z_{ds}^{R} P_R + \delta Z_{ds}^{L} m_s P_R - m_d \delta Z_{ds}^{L} P_L \Big\} u_s(p_1) \qquad \text{(D.17)}
\end{aligned}
$$

$$
\begin{aligned}
=& \frac{-i(p_1 - p_2)_\mu}{2} \, i g_d \, \bar{u}_d(p_2) \gamma^\mu P_R \frac{1}{m_s^2 - m_d^2} \\
& \Big\{ m_d \Big(\delta Z_{ds}^{R} m_s P_L - m_d \delta Z_{ds}^{R} P_R + \delta Z_{ds}^{L} m_s P_R - m_d \delta Z_{ds}^{L} P_L \Big) \\
& + m_s \Big(\delta Z_{ds}^{R} m_s P_R - m_d \delta Z_{ds}^{R} P_L + \delta Z_{ds}^{L} m_s P_L - m_d \delta Z_{ds}^{L} P_R \Big) \Big\} u_s(p_1) \\
& \qquad \text{(D.18)} \\
=& \frac{-i(p_1 - p_2)_\mu}{2} \, i g_d \, \bar{u}_d(p_2) \gamma^\mu \, \delta Z_{ds}^{R} \, P_R u_s(p_1) \, . \qquad \text{(D.19)}
\end{aligned}
$$

The result is the negative g_d term of equation (D.10). Thus, adding all three diagrams together gives

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

independently of the precise form of the renormalization constants.

E Primakoff form factor beyond the static limit

This appendix is based on results and arguments which were published in appendix A of ref. [3] before submission of this thesis. Material from this article is used in the following without further reference.

The derivations of the Primakoff rates Γ_a^{P} in equations (7.2) and (7.9) assume a form factor for the scattering potential that is given by [231]

$$|F_{\mathrm{eff}}|^2 = Z^2 \frac{q^2}{q^2 + \kappa_{\mathrm{s}}^2}\,, \tag{E.1}$$

where Z is the charge in units of the elementary charge, q is the magnitude of the transferred momentum $\vec{q}$ and κ_{s} is the usual screening scale.

This form factor was derived in ref. [231] by applying the strict static limit, hence assuming that the scattering potential does not change during the time it takes for one plasmon to cross the potential. Since charges are screened at typical distances of $1/\kappa_{\mathrm{s}}$, the corresponding time scale is given by $t \sim 1/\kappa_{\mathrm{s}}$.

The opposing limit would be to assume a Yukawa potential for every scattering plasmon. Since the results of these two limiting cases differ by a large factor ~ 100, it is necessary to investigate the size of the expected corrections once the assumption of static point charges is dropped. To do this, the calculation in the static limit as presented in ref. [231] is briefly recapped and then modified to include point charges moving at constant velocity.

Static limit

It is helpful to recall the standard result for a situation where the plasmon encounters a set of non-moving charges Z_i. This was first discussed in ref. [231].

The form factor of N particles with charges Z_i at positions $\vec{r}_i$ is given by

$$F_N(\vec{q}) = \sum_{i=1}^{N} Z_i \, e^{i\,\vec{q}\cdot\vec{r}_i}\,, \tag{E.2}$$

and hence

$$|F_N(\vec{q})|^2 = \sum_i Z_i^2 + \sum_{\substack{i,j\\ i\neq j}} Z_i Z_j \cos(\vec{q}\cdot\vec{r}_{ij})\,, \tag{E.3}$$

where $\vec{r}_{ij}$ are the inter-particle distances. In order to obtain the effective form factor in a plasma, this has to be averaged over the locations of charges. This requires the probability density of the distance between two point charges [231, 430],

$$p_{ij} = \frac{1}{V}\left(e^{-\Lambda r} - \frac{Z_i Z_j \alpha_{\mathrm{EM}}}{T}\frac{e^{-\kappa_{\mathrm{s}} r}}{r}\right) , \tag{E.4}$$

where r is the distance between particles i and j. To simplify the calculation of the resulting volume integrals, it is helpful to include a regulator Λ, which will be sent to 0 in the end. V is a volume normalization factor that converges to the full volume in the limit of $\Lambda \to 0$. Averaging $|F_N|^2$ over the inter-particle distances yields

$$\begin{aligned} \langle |F_N(\vec{q})|^2\rangle_r =& \sum_i Z_i^2 + \sum_{\substack{i,j\\ i\neq j}} Z_i Z_j \int \mathrm{d}r^3\, p_{ij}(r)\cos(\vec{q}\cdot\vec{r}) \\ =& \sum_i Z_i^2 + 4\pi \sum_{\substack{i,j\\ i\neq j}} Z_i Z_j \int_0^\infty \mathrm{d}r\, r^2 \frac{1}{V}\left(e^{-\Lambda r} - \frac{Z_i Z_j \alpha_{\mathrm{EM}}}{T}\frac{e^{-\kappa_{\mathrm{s}} r}}{r}\right)\frac{\sin(q\,r)}{q\,r} . \end{aligned} \tag{E.5}$$

The first term in parentheses vanishes in the limit $\Lambda \to 0$, while the second one gives a finite contribution. Using that the screening scale, as in equation (7.4), can be written as

$$\kappa_{\mathrm{s}}^2 = \frac{4\pi\alpha_{\mathrm{EM}}}{T}\sum_i \frac{Z_i^2}{V} , \tag{E.6}$$

the sum becomes

$$\sum_{\substack{i,j\\ i\neq j}} Z_i^2 Z_j^2 = \sum_i Z_i^2 \sum_{j\neq i} Z_j^2 \simeq \sum_i Z_i^2 \frac{\kappa_{\mathrm{s}}^2 T V}{4\pi\alpha_{\mathrm{EM}}} , \tag{E.7}$$

where the last equality is exact in the large volume limit. By putting everything together, the result of ref. [231] is recovered,

$$\langle |F_N(\vec{q})|^2\rangle_r = \sum_i Z_i^2\left(1 - \frac{\kappa_{\mathrm{s}}^2}{\kappa_{\mathrm{s}}^2 + q^2}\right) = \sum_i Z_i^2 \frac{q^2}{\kappa_{\mathrm{s}}^2 + q^2} . \tag{E.8}$$

Beyond the static limit

To see how decoherence can affect the expected form factor, the strict static limit is now dropped and instead it is assumed that each particle i moves with a constant velocity $\vec{v}_i$. As a result, the form factor is now time-dependent,

$$F_N(\vec{q}) = \sum_i Z_i\, e^{i\,\vec{q}\cdot(\vec{r}+\vec{v}_i t)} , \tag{E.9}$$

and it can be averaged over the time $t \sim 1/\kappa_{\mathrm{s}}$ of one scattering event,

$$\langle F_N(\vec{q})\rangle_t = \frac{1}{t}\int_{-\frac{t}{2}}^{\frac{t}{2}} \mathrm{d}t' \sum_i Z_i\, e^{i\,\vec{q}\cdot(\vec{r}+\vec{v}_i t')} = \sum_i Z_i\, \mathrm{sinc}(\Delta\phi_i)\, e^{i\,\vec{q}\cdot\vec{r}} , \tag{E.10}$$

where $\Delta\phi_i \equiv \frac{1}{2}\vec{q}\cdot\vec{v}_i t$ denotes the phase shift of particle i.

The average over the inter-particle distance can now be evaluated exactly as before – up to the point of equation (E.5) – by just carrying along the factors of $\text{sinc}(\Delta\phi_i)$,

$$\begin{aligned}\langle|\langle F_N(\vec{q})\rangle_t|^2\rangle_r = \sum_i Z_i^2\,\text{sinc}^2(\Delta\phi_i) + \frac{4\pi}{V}\sum_{\substack{i,j\\ i\neq j}} Z_i Z_j \text{sinc}(\Delta\phi_i)\text{sinc}(\Delta\phi_j)\\ \times\int_0^\infty \text{d}r\,r^2\left(e^{-\Lambda r} - \frac{Z_i Z_j \alpha_{\text{EM}}}{T}\frac{e^{-\kappa_\text{s} r}}{r}\right)\frac{\sin(q\,r)}{q\,r}\,. \qquad \text{(E.11)}\end{aligned}$$

As before, the first term in parentheses vanishes in limit $\Lambda\to 0$ while the second term is finite. The sum over pairs of charges is modified by the suppression factors.

$$\begin{aligned}\sum_{\substack{i,j\\ i\neq j}}\text{sinc}(\Delta\phi_i)Z_i^2\text{sinc}(\Delta\phi_j)Z_j^2 &= \sum_i \text{sinc}(\Delta\phi_i)Z_i^2\sum_{\substack{j\\ j\neq i}}\text{sinc}(\Delta\phi_j)Z_j^2 && \text{(E.12)}\\ &\simeq \sum_i \text{sinc}(\Delta\phi_i)Z_i^2\frac{\kappa_\text{eff}^2 TV}{4\pi\alpha_\text{EM}}\,, && \text{(E.13)}\end{aligned}$$

where in the last line, κ_eff^2 was defined as

$$\kappa_\text{eff}^2 \equiv \frac{4\pi\alpha_\text{EM}}{T}\sum_i\frac{\text{sinc}(\Delta\phi_i)Z_i^2}{V}. \qquad \text{(E.14)}$$

Inserting these results into the expression for $\langle|\langle F_N(\vec{q})\rangle_t|^2\rangle_r$, one arrives at

$$\begin{aligned}\langle|\langle F_N(\vec{q})\rangle_t|^2\rangle_r &= \sum_i \text{sinc}(\Delta\phi_i)Z_i^2\left(\text{sinc}(\Delta\phi_i) - \kappa_\text{eff}^2\int_0^\infty \text{d}r\,e^{-\kappa_\text{s} r}\frac{\sin(q\,r)}{q}\right) && \text{(E.15)}\\ &= \sum_i Z_i^2\,\text{sinc}(\Delta\phi_i)\left(\text{sinc}(\Delta\phi_i) - \frac{\kappa_\text{eff}^2}{\kappa_\text{s}^2+q^2}\right) && \text{(E.16)}\\ &= \sum_i Z_i^2\,\text{sinc}(\Delta\phi_i)\left(\frac{\text{sinc}(\Delta\phi_i)\kappa_\text{s}^2 - \kappa_\text{eff}^2 + \text{sinc}(\Delta\phi_i)q^2}{\kappa_\text{s}^2+q^2}\right)\,. && \text{(E.17)}\end{aligned}$$

As expected, this result converges to the previous one in the limit $\Delta\phi_i\to 0$, where $\kappa_\text{eff}\to\kappa_\text{s}$.

To evaluate the relative correction of the new result compared to the static limit, the thermal averages of the suppression factors $\text{sinc}(\Delta\phi_i)$ and $\text{sinc}^2(\Delta\phi_i)$ need to be evaluated. For simplicity, one can assume a Maxwell-Boltzmann distribution for electrons and use typical values for the solar core temperature. By expanding the sinc function, the estimated values can be given as

$$\begin{aligned}\langle\text{sinc}(\Delta\phi_i)\rangle_{v_i}\big|_e &= \int\text{d}^3v\,\text{sinc}(\Delta\phi_i)\left(\frac{m_e}{2\pi T}\right)^{\frac{3}{2}}e^{-\frac{m_e v^2}{2T}} \sim 1-\left(10^{-4}\,\frac{q^2}{\kappa_\text{s}^2}\right) && \text{(E.18)}\\ \langle\text{sinc}^2(\Delta\phi_i)\rangle_{v_i}\big|_e &= \int\text{d}^3v\,\text{sinc}^2(\Delta\phi_i)\left(\frac{m_e}{2\pi T}\right)^{\frac{3}{2}}e^{-\frac{m_e v^2}{2T}} \sim 1-\left(2\times 10^{-4}\,\frac{q^2}{\kappa_\text{s}^2}\right)\,. && \text{(E.19)}\end{aligned}$$

These factors for electrons are already very close to unity and it is safe to neglect the equivalent suppression factors for significantly heavier ions.

The effective screening scale as defined in equation (E.14) still depends on the velocities of particles. But using equation (E.18) and the fact that electrons contribute at most 45 % to the square of the screening scale (as shown in figure 7.2), one can see that the thermal average of the difference between κ_s^2 and κ_eff^2 is given by

$$\kappa_\mathrm{s}^2 - \langle \kappa_\mathrm{eff}^2 \rangle_{v_i} \lesssim 0.45 \times 10^{-4} \, q^2 \, . \tag{E.20}$$

This means that the resulting relative correction to the square of the form factor in equation (E.17) does not exceed 5×10^{-5}. Furthermore, equations (E.18) and (E.19) indicate that the sinc factors in equation (E.17) result in corrections by at most 10^{-4} when it is conservatively assumed that $q \lesssim \kappa_\mathrm{s}$.

In conclusion, it is expected that the square of the form factor, and thereby the total Primakoff rate, only changes by $\lesssim 0.02\,\%$ when the assumption of a static distribution of point-like charges is dropped.

F Axions from other nuclei than ^{57}Fe

This appendix is based on results and arguments which were published in appendix A of ref. [4] before submission of this thesis. Material from this article is used in the following without further reference.

The phenomenological discussion in chapter 7 centres around the 14.4 keV line of ^{57}Fe because it is expected to give the strongest signal. In order to ensure that this is true, especially considering that the ^{57}Fe-line suffers from strong thermal suppression, this appendix systematically searches for alternative nuclear M1 transitions which may also generate a line in the solar axion flux.

A list of potential candidates is provided in the appendix of ref. [373] in form of a list of isotopes featuring low-energy nuclear transitions. All of the calculations in section 7.1.5 equally apply to M1 transitions of nuclei other than ^{57}Fe. Hence, one only needs to compare their respective axion flux per solar mass. This is expressed in equation (7.47) as

$$\mathcal{N}_a = \mathcal{N}\omega_1 \frac{1}{\tau_0}\frac{1}{1+\alpha}\frac{\Gamma_a}{\Gamma_\gamma} \ . \tag{F.1}$$

The nuclear matrix elements entering in the ratio of axion to photon emissions have not been computed with equal precision for the various nuclear transitions. It is however reasonable to assume values of $\mathcal{O}(1)$ for the dimensionless constants β and η. Furthermore, the E2/M1 mixing ratio is already close to its ideal value of zero for the ^{57}Fe line. The factor $\frac{\Gamma_a}{\Gamma_\gamma}$ is therefore expected to be comparable (or smaller than the one of ^{57}Fe in case of a large δ) for all nuclear transitions and it is sufficient to focus on the combination of abundance, occupation number and inverse lifetime.

Very large differences appear in the respective number density $\mathcal{N}$ of isotopes. To estimate their value for all radii in the Sun, two assumptions have to be made. First that the contribution of one isotope to the total element abundance is constant throughout the Sun and identical to the one found on earth, which is denoted as a. And second that the radial density profiles of heavy elements only differ from the one of iron by a constant factor. These two reasonable assumptions allow to compare the respective values of $\mathcal{N}$ by just multiplying the photospheric abundance with the relative isotope abundance on earth. The photospheric abundance is tabulated in ref. [374] on a logarithmic scale. The ratio of the number density of the element in question normalized to the one of hydrogen, i.e. $\epsilon \equiv N_X/N_H$, is given in table F.1.

The thermal occupation number ω_1 crucially depends on the transition energy E^* and can be written as [203, 370]

$$\omega_1 = \frac{(2J_1+1)e^{-E^*/T}}{(2J_0+1)+(2J_1+1)e^{-E^*/T}}, \tag{F.2}$$

	^{57}Fe	**^{83}Kr**	**^{169}Tm**	**^{187}Os**	**^{201}Hg**
E^* [keV]	14.4	9.4	8.4	9.7	1.6
J_0	1/2	9/2	1/2	1/2	3/2
J_1	3/2	7/2	3/2	3/2	1/2
τ_0 [ns]	141	212	5.9	3.4	144
α	8.56	17.09	285	264	47000
ϵ	$10^{-4.5}$	$10^{-8.75}$	$10^{-11.9}$	$10^{-10.6}$	$10^{-10.83}$
a [%]	2.14	11.55	100	1.6	13.2
$\mathcal{N}_a(r=0)$ [relative to ^{57}Fe]	1	1.8×10^{-3}	1.3×10^{-4}	3.0×10^{-5}	1.9×10^{-6}

Table F.1. Isotopes with a nuclear M1 transition and $E^* < 20\,\text{keV}$. The element abundances ϵ are taken from ref. [374]. All other values are tabled in the appendix of ref. [373]. The values in the last row were calculated by evaluating equations (F.1) and (F.2) with the solar core temperature $T(r=0) = 1.33\,\text{keV}$.

where J_0 and J_1 are the total angular momentum quantum numbers of the ground and excited states, respectively. Because of the exponential suppression, the transition energy E^* has a large effect on the resulting axion flux.

All isotopes from the appendix of ref. [373] with a nuclear M1 transition below 20 keV were collected in table F.1. The last row indicates the size of the axion flux from these nuclear transitions relative to the one from ^{57}Fe. It becomes clear that the strong Boltzmann suppression in the case of ^{57}Fe is mitigated by a relatively large abundance and a small internal conversion coefficient. Even though the values in table F.1 should be understood as very rough estimates, this clearly indicates that the ^{57}Fe line generates the strongest axion signal from nuclear deexcitations.

Bibliography

[1] G. Alonso-Álvarez, F. Ertas, J. Jaeckel, F. Kahlhoefer and L. J. Thormaehlen, *Hidden Photon Dark Matter in the Light of XENON1T and Stellar Cooling*, *JCAP* **11** (2020) 029, [`2006.11243`].

[2] G. Alonso-Álvarez, F. Ertas, J. Jaeckel, F. Kahlhoefer and L. J. Thormaehlen, *Leading logs in QCD axion effective field theory*, *JHEP* **07** (2021) 059, [`2101.03173`].

[3] S. Hoof, J. Jaeckel and L. J. Thormaehlen, *Quantifying uncertainties in the solar axion flux and their impact on determining axion model parameters*, *JCAP* **09** (2021) 006, [`2101.08789`].

[4] L. Di Luzio et al., *Probing the axion–nucleon coupling with the next generation of axion helioscopes*, *Eur. Phys. J. C* **82** (2022) 120, [`2111.06407`].

[5] M. D. Schwartz, *Quantum Field Theory and the Standard Model.* Cambridge University Press, 2014.

[6] S. L. Glashow, *Partial Symmetries of Weak Interactions*, *Nucl. Phys.* **22** (1961) 579–588.

[7] A. Salam and J. C. Ward, *Electromagnetic and weak interactions*, *Phys. Lett.* **13** (1964) 168–171.

[8] S. Weinberg, *A Model of Leptons*, *Phys. Rev. Lett.* **19** (1967) 1264–1266.

[9] A. Salam, *Weak and Electromagnetic Interactions*, *Conf. Proc. C* **680519** (1968) 367–377.

[10] CMS collaboration, S. Chatrchyan et al., *Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC*, *Phys. Lett. B* **716** (2012) 30–61, [`1207.7235`].

[11] ATLAS collaboration, G. Aad et al., *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, *Phys. Lett. B* **716** (2012) 1–29, [`1207.7214`].

[12] F. Wilczek, *Problem of Strong P and T Invariance in the Presence of Instantons*, *Phys. Rev. Lett.* **40** (1978) 279–282.

[13] S. Weinberg, *A New Light Boson?*, *Phys. Rev. Lett.* **40** (1978) 223–226.

[14] R. D. Peccei and H. R. Quinn, *CP Conservation in the Presence of Pseudoparticles*, *Phys. Rev. Lett.* **38** (1977) 1440–1443.

[15] J. Preskill, M. B. Wise and F. Wilczek, *Cosmology of the Invisible Axion*, *Phys. Lett. B* **120** (1983) 127–132.

[16] L. F. Abbott and P. Sikivie, *A Cosmological Bound on the Invisible Axion*, *Phys. Lett. B* **120** (1983) 133–136.

[17] G. Servant, *Baryogenesis from Strong CP Violation and the QCD Axion*, *Phys. Rev. Lett.* **113** (2014) 171803, [`1407.0030`].

[18] A. D. Linde, *Axions in inflationary cosmology*, *Phys. Lett. B* **259** (1991) 38–47.

[19] K. Freese, J. T. Liu and D. Spolyar, *Inflating with the QCD axion*, *Phys. Rev. D* **72** (2005) 123521, [`hep-ph/0502177`].

[20] T. W. Grimm, *Axion inflation in type II string theory*, *Phys. Rev. D* **77** (2008) 126007, [`0710.3883`].

[21] L. McAllister, E. Silverstein and A. Westphal, *Gravity Waves and Linear Inflation from Axion Monodromy*, *Phys. Rev. D* **82** (2010) 046003, [`0808.0706`].

[22] F. Marchesano, G. Shiu and A. M. Uranga, *F-term Axion Monodromy Inflation*, *JHEP* **09** (2014) 184, [`1404.3040`].

[23] Y. Ema, K. Hamaguchi, T. Moroi and K. Nakayama, *Flaxion: a minimal extension to solve puzzles in the standard model*, *JHEP* **01** (2017) 096, [`1612.05492`].

[24] P. Langacker, R. D. Peccei and T. Yanagida, *Invisible Axions and Light Neutrinos: Are They Connected?*, *Mod. Phys. Lett. A* **1** (1986) 541.

[25] M. Shin, *Light Neutrino Masses and Strong CP Problem*, *Phys. Rev. Lett.* **59** (1987) 2515.

[26] A. G. Dias, A. C. B. Machado, C. C. Nishi, A. Ringwald and P. Vaudrevange, *The Quest for an Intermediate-Scale Accidental Axion and Further ALPs*, *JHEP* **06** (2014) 037, [`1403.5760`].

[27] M. Kamionkowski, J. Pradler and D. G. E. Walker, *Dark energy from the string axiverse*, *Phys. Rev. Lett.* **113** (2014) 251302, [`1409.0549`].

[28] J. Jaeckel and A. Ringwald, *The Low-Energy Frontier of Particle Physics*, *Ann. Rev. Nucl. Part. Sci.* **60** (2010) 405–437, [`1002.0329`].

[29] J. E. Kim, *Light Pseudoscalars, Particle Physics and Cosmology*, *Phys. Rept.* **150** (1987) 1–177.

[30] J. E. Kim and G. Carosi, *Axions and the Strong CP Problem*, *Rev. Mod. Phys.* **82** (2010) 557–602, [`0807.3125`].

[31] S. Vandoren and P. van Nieuwenhuizen, *Lectures on instantons*, [`0802.1862`].

[32] A. A. Belavin, A. M. Polyakov, A. S. Schwartz and Y. S. Tyupkin, *Pseudoparticle Solutions of the Yang-Mills Equations*, *Phys. Lett. B* **59** (1975) 85–87.

[33] K. M. Bitar and S.-J. Chang, *Vacuum Tunneling of Gauge Theory in Minkowski Space*, *Phys. Rev. D* **17** (1978) 486.

[34] C. G. Callan, Jr., R. F. Dashen and D. J. Gross, *Toward a Theory of the Strong Interactions*, *Phys. Rev. D* **17** (1978) 2717.

[35] G. 't Hooft, *Computation of the Quantum Effects Due to a Four-Dimensional Pseudoparticle*, *Phys. Rev. D* **14** (1976) 3432–3450.

[36] M. H. G. Tytgat, *QCD at theta similar to pi reexamined: Domain walls and spontaneous CP violation*, *Phys. Rev. D* **61** (2000) 114009, [`hep-ph/9909532`].

[37] C. Abel et al., *Measurement of the Permanent Electric Dipole Moment of the Neutron*, *Phys. Rev. Lett.* **124** (2020) 081803, [`2001.11966`].

[38] R. Alarcon et al., *Electric dipole moments and the search for new physics*, in *2022 Snowmass Summer Study*, 3, 2022, [`2203.08103`].

[39] M. Pospelov and A. Ritz, *Theta vacua, QCD sum rules, and the neutron electric dipole moment*, *Nucl. Phys. B* **573** (2000) 177–200, [`hep-ph/9908508`].

[40] V. Baluni, *CP Violating Effects in QCD*, *Phys. Rev. D* **19** (1979) 2227–2230.

[41] A. Pich and E. de Rafael, *Strong CP violation in an effective chiral Lagrangian approach*, *Nucl. Phys. B* **367** (1991) 313–333.

[42] F. K. Guo, R. Horsley, U. G. Meissner, Y. Nakamura, H. Perlt, P. E. L. Rakow et al., *The electric dipole moment of the neutron from 2+1 flavor lattice QCD*, *Phys. Rev. Lett.* **115** (2015) 062001, [`1502.02295`].

[43] E. Shintani, T. Blum, T. Izubuchi and A. Soni, *Neutron and proton electric dipole moments from* $N_f = 2 + 1$ *domain-wall fermion lattice QCD*, *Phys. Rev. D* **93** (2016) 094503, [`1512.00566`].

[44] M. Abramczyk, S. Aoki, T. Blum, T. Izubuchi, H. Ohki and S. Syritsyn, *Lattice calculation of electric dipole moments and form factors of the nucleon*, *Phys. Rev. D* **96** (2017) 014501, [`1701.07792`].

[45] J. Dragos, T. Luu, A. Shindler, J. de Vries and A. Yousif, *Confirming the Existence of the strong CP Problem in Lattice QCD with the Gradient Flow*, *Phys. Rev. C* **103** (2021) 015202, [`1902.03254`].

[46] C. Vafa and E. Witten, *Parity Conservation in QCD*, *Phys. Rev. Lett.* **53** (1984) 535.

[47] G. 't Hooft, C. Itzykson, A. Jaffe, H. Lehmann, P. K. Mitter, I. M. Singer et al., *Recent Developments in Gauge Theories. Proceedings, Nato Advanced Study Institute, Cargese, France, August 26 - September 8, 1979*, *NATO Sci. Ser. B* **59** (1980) pp. 1–438.

[48] Particle Data Group collaboration, M. Tanabashi et al., *Review of Particle Physics*, *Phys. Rev. D* **98** (2018) 030001.

[49] S. L. Adler, *Axial vector vertex in spinor electrodynamics*, *Phys. Rev.* **177** (1969) 2426–2438.

[50] J. S. Bell and R. Jackiw, *A PCAC puzzle: $\pi^0 \to \gamma\gamma$ in the σ model*, *Nuovo Cim. A* **60** (1969) 47–61.

[51] RM123 collaboration, G. M. de Divitiis, R. Frezzotti, V. Lubicz, G. Martinelli, R. Petronzio, G. C. Rossi et al., *Leading isospin breaking effects on the lattice*, *Phys. Rev. D* **87** (2013) 114505, [`1303.4896`].

[52] MILC collaboration, S. Basak et al., *Electromagnetic effects on the light hadron spectrum*, *J. Phys. Conf. Ser.* **640** (2015) 012052, [`1510.04997`].

[53] R. Horsley et al., *Isospin splittings of meson and baryon masses from three-flavor lattice QCD + QED*, *J. Phys. G* **43** (2016) 10LT02, [`1508.06401`].

[54] G. Grilli di Cortona, E. Hardy, J. Pardo Vega and G. Villadoro, *The QCD axion, precisely*, *JHEP* **01** (2016) 034, [`1511.02867`].

[55] A. A. Anselm and A. A. Johansen, *Can electroweak theta term be observable?*, *Nucl. Phys. B* **412** (1994) 553–573, [`hep-ph/9305271`].

[56] P. Fileviez Perez and H. H. Patel, *The Electroweak Vacuum Angle*, *Phys. Lett. B* **732** (2014) 241–243, [`1402.6340`].

[57] D. Tong, *Line Operators in the Standard Model*, *JHEP* **07** (2017) 104, [`1705.01853`].

[58] R. D. Peccei and H. R. Quinn, *Constraints Imposed by CP Conservation in the Presence of Pseudoparticles*, *Phys. Rev. D* **16** (1977) 1791–1797.

[59] F. Wilczek, *The Birth of Axions*, *Current Contents* (1991) 8–9.

[60] G. Alexanian, R. MacKenzie, M. B. Paranjape and J. Ruel, *Path integration and perturbation theory with complex Euclidean actions*, *Phys. Rev. D* **77** (2008) 105014, [`0802.0354`].

[61] M. Kobayashi and T. Maskawa, *CP Violation in the Renormalizable Theory of Weak Interaction*, *Prog. Theor. Phys.* **49** (1973) 652–657.

[62] S. Dar, *The Neutron EDM in the SM: A Review*, [`hep-ph/0008248`].

[63] M. Pospelov and A. Ritz, *Electric dipole moments as probes of new physics*, *Annals Phys.* **318** (2005) 119–169, [`hep-ph/0504231`].

[64] J. Beacham et al., *Physics Beyond Colliders at CERN: Beyond the Standard Model Working Group Report*, [`1901.09966`].

[65] P. Agrawal and K. Howe, *Factoring the Strong CP Problem*, *JHEP* **12** (2018) 029, [`1710.04213`].

[66] D. S. M. Alves and N. Weiner, *A viable QCD axion in the MeV mass range*, *JHEP* **07** (2018) 092, [`1710.03764`].

[67] 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, [`1805.06465`].

[68] T. Gherghetta, V. V. Khoze, A. Pomarol and Y. Shirman, *The Axion Mass from 5D Small Instantons*, *JHEP* **03** (2020) 063, [`2001.05610`].

[69] M. Gorghetto and G. Villadoro, *Topological Susceptibility and QCD Axion Mass: QED and NNLO corrections*, *JHEP* **03** (2019) 033, [`1812.01008`].

[70] European Twisted Mass collaboration, N. Carrasco et al., *Up, down, strange and charm quark masses with $N_f = 2+1+1$ twisted mass lattice QCD*, *Nucl. Phys. B* **887** (2014) 19–68, [`1403.4504`].

[71] Z. Fodor, C. Hoelbling, S. Krieg, L. Lellouch, T. Lippert, A. Portelli et al., *Up and down quark masses and corrections to Dashen's theorem from lattice QCD and quenched QED*, *Phys. Rev. Lett.* **117** (2016) 082001, [`1604.07112`].

[72] D. Giusti, V. Lubicz, C. Tarantino, G. Martinelli, F. Sanfilippo, S. Simula et al., *Leading isospin-breaking corrections to pion, kaon and charmed-meson masses with Twisted-Mass fermions*, *Phys. Rev. D* **95** (2017) 114504, [`1704.06561`].

[73] MILC collaboration, S. Basak et al., *Lattice computation of the electromagnetic contributions to kaon and pion masses*, *Phys. Rev. D* **99** (2019) 034503, [`1807.05556`].

[74] S. Dodelson, *Modern Cosmology.* Academic Press, Amsterdam, 2003.

[75] K. Freese, *Review of Observational Evidence for Dark Matter in the Universe and in upcoming searches for Dark Stars*, *EAS Publ. Ser.* **36** (2009) 113–126, [`0812.4005`].

[76] M. Bauer and T. Plehn, *Yet Another Introduction to Dark Matter: The Particle Physics Approach*, vol. 959 of *Lecture Notes in Physics.* Springer, 2019, 10.1007/978-3-030-16234-4.

[77] A. Arbey and F. Mahmoudi, *Dark matter and the early Universe: a review*, *Prog. Part. Nucl. Phys.* **119** (2021) 103865, [`2104.11488`].

[78] Planck collaboration, N. Aghanim et al., *Planck 2018 results. I. Overview and the cosmological legacy of Planck*, *Astron. Astrophys.* **641** (2020) A1, [`1807.06205`].

[79] V. C. Rubin and J. Ford, W. Kent, *Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions*, *Astrophys. J.* **159** (Feb., 1970) 379.

[80] M. Bartelmann, *Gravitational Lensing*, *Class. Quant. Grav.* **27** (2010) 233001, [`1010.3829`].

[81] D. Clowe, M. Bradac, A. H. Gonzalez, M. Markevitch, S. W. Randall, C. Jones et al., *A direct empirical proof of the existence of dark matter*, *Astrophys. J. Lett.* **648** (2006) L109–L113, [`astro-ph/0608407`].

[82] R. A. Flores and J. R. Primack, *Observational and theoretical constraints on singular dark matter halos*, *Astrophys. J. Lett.* **427** (1994) L1–4, [`astro-ph/9402004`].

[83] B. Moore, *Evidence against dissipation-less dark matter from observations of galaxy haloes*, *Nature* **370** (Aug., 1994) 629–631.

[84] A. A. Klypin, A. V. Kravtsov, O. Valenzuela and F. Prada, *Where are the missing Galactic satellites?*, *Astrophys. J.* **522** (1999) 82–92, [`astro-ph/9901240`].

[85] B. Moore, S. Ghigna, F. Governato, G. Lake, T. R. Quinn, J. Stadel et al., *Dark matter substructure within galactic halos*, *Astrophys. J. Lett.* **524** (1999) L19–L22, [`astro-ph/9907411`].

[86] M. Boylan-Kolchin, J. S. Bullock and M. Kaplinghat, *Too big to fail? The puzzling darkness of massive Milky Way subhaloes*, *MNRAS* **415** (July, 2011) L40–L44, [`1103.0007`].

[87] T. Sawala et al., *The APOSTLE simulations: solutions to the Local Group's cosmic puzzles*, *Mon. Not. Roy. Astron. Soc.* **457** (2016) 1931–1943, [`1511.01098`].

[88] A. Benitez-Llambay and C. Frenk, *The detailed structure and the onset of galaxy formation in low-mass gaseous dark matter haloes*, *MNRAS* **498** (Nov., 2020) 4887–4900, [`2004.06124`].

[89] Q. Gu, Q. Guo, T. Zhang, M. Cautun, C. Lacey, C. S. Frenk et al., *The spatial distribution of satellites in galaxy clusters*, *MNRAS* **514** (May, 2022) 390–402, [`2205.06767`].

[90] B. Carr, F. Kuhnel and M. Sandstad, *Primordial Black Holes as Dark Matter*, *Phys. Rev. D* **94** (2016) 083504, [`1607.06077`].

[91] A. M. Green and B. J. Kavanagh, *Primordial Black Holes as a dark matter candidate*, *J. Phys. G* **48** (2021) 043001, [`2007.10722`].

[92] S. Tremaine and J. E. Gunn, *Dynamical Role of Light Neutral Leptons in Cosmology*, *Phys. Rev. Lett.* **42** (1979) 407–410.

[93] P. Arias, D. Cadamuro, M. Goodsell, J. Jaeckel, J. Redondo and A. Ringwald, *WISPy Cold Dark Matter*, *JCAP* **06** (2012) 013, [`1201.5902`].

[94] D. J. E. Marsh, *Axion Cosmology*, *Phys. Rept.* **643** (2016) 1–79, [`1510.07633`].

[95] P. Fox, A. Pierce and S. D. Thomas, *Probing a QCD string axion with precision cosmological measurements*, [`hep-th/0409059`].

[96] V. B. . Klaer and G. D. Moore, *The dark-matter axion mass*, *JCAP* **11** (2017) 049, [`1708.07521`].

[97] M. Gorghetto, E. Hardy and G. Villadoro, *Axions from Strings: the Attractive Solution*, *JHEP* **07** (2018) 151, [`1806.04677`].

[98] M. Hindmarsh, J. Lizarraga, A. Lopez-Eiguren and J. Urrestilla, *Scaling Density of Axion Strings*, *Phys. Rev. Lett.* **124** (2020) 021301, [`1908.03522`].

[99] M. Gorghetto, E. Hardy and G. Villadoro, *More axions from strings*, *SciPost Phys.* **10** (2021) 050, [`2007.04990`].

[100] M. Buschmann, J. W. Foster, A. Hook, A. Peterson, D. E. Willcox, W. Zhang et al., *Dark matter from axion strings with adaptive mesh refinement*, *Nature Commun.* **13** (2022) 1049, [`2108.05368`].

[101] L. Caloni, M. Gerbino, M. Lattanzi and L. Visinelli, *Novel cosmological bounds on thermally-produced axion-like particles*, [`2205.01637`].

[102] I. J. Allali, M. P. Hertzberg and Y. Lyu, *Altered Axion Abundance from a Dynamical Peccei-Quinn Scale*, *Phys. Rev. D* **105** (2022) 123517, [`2203.15817`].

[103] A. Papageorgiou, P. Quílez and K. Schmitz, *Axion dark matter from frictional misalignment*, [`2206.01129`].

[104] K. Choi, S. H. Im, H. J. Kim and H. Seong, *Axion dark matter with thermal friction*, [`2206.01462`].

[105] J. Redondo, *Solar axion flux from the axion-electron coupling*, *JCAP* **1312** (2013) 008, [`1310.0823`].

[106] G. Alonso-Álvarez, M. B. Gavela and P. Quilez, *Axion couplings to electroweak gauge bosons*, *Eur. Phys. J. C* **79** (2019) 223, [`1811.05466`].

[107] W. A. Bardeen, *Anomalous Ward identities in spinor field theories*, *Phys. Rev.* **184** (1969) 1848–1857.

[108] L. Di Luzio, F. Mescia and E. Nardi, *Redefining the Axion Window*, *Phys. Rev. Lett.* **118** (2017) 031801, [`1610.07593`].

[109] J. E. Kim, *Weak Interaction Singlet and Strong CP Invariance*, *Phys. Rev. Lett.* **43** (1979) 103.

[110] M. A. Shifman, A. I. Vainshtein and V. I. Zakharov, *Can Confinement Ensure Natural CP Invariance of Strong Interactions?*, *Nucl. Phys. B* **166** (1980) 493–506.

[111] E. Nardi and E. Roulet, *Are exotic stable quarks cosmologically allowed?*, *Phys. Lett. B* **245** (1990) 105–110.

[112] M. L. Perl, P. C. Kim, V. Halyo, E. R. Lee, I. T. Lee, D. Loomba et al., *The Search for stable, massive, elementary particles*, *Int. J. Mod. Phys. A* **16** (2001) 2137–2164, [`hep-ex/0102033`].

[113] M. L. Perl, E. R. Lee and D. Loomba, *A Brief review of the search for isolatable fractional charge elementary particles*, *Mod. Phys. Lett. A* **19** (2004) 2595–2610.

[114] J. Quevillon and C. Smith, *Axions are blind to anomalies*, *Eur. Phys. J. C* **79** (2019) 822, [`1903.12559`].

[115] L. Di Luzio, F. Mescia and E. Nardi, *Window for preferred axion models*, *Phys. Rev. D* **96** (2017) 075003, [`1705.05370`].

[116] A. R. Zhitnitsky, *On Possible Suppression of the Axion Hadron Interactions. (In Russian)*, *Sov. J. Nucl. Phys.* **31** (1980) 260.

[117] M. Dine, W. Fischler and M. Srednicki, *A Simple Solution to the Strong CP Problem with a Harmless Axion*, *Phys. Lett. B* **104** (1981) 199–202.

[118] I. G. Irastorza and J. Redondo, *New experimental approaches in the search for axion-like particles*, *Prog. Part. Nucl. Phys.* **102** (2018) 89–159, [`1801.08127`].

[119] F. Björkeroth, L. Di Luzio, F. Mescia, E. Nardi, P. Panci and R. Ziegler, *Axion-electron decoupling in nucleophobic axion models*, *Phys. Rev. D* **101** (2020) 035027, [`1907.06575`].

[120] A. G. Dias, J. Leite, J. W. F. Valle and C. A. Vaquera-Araujo, *Reloading the axion in a 3-3-1 setup*, *Phys. Lett. B* **810** (2020) 135829, [`2008.10650`].

[121] L. Di Luzio, F. Mescia, E. Nardi, P. Panci and R. Ziegler, *Astrophobic Axions*, *Phys. Rev. Lett.* **120** (2018) 261803, [`1712.04940`].

[122] F. Björkeroth, L. Di Luzio, F. Mescia and E. Nardi, *$U(1)$ flavour symmetries as Peccei-Quinn symmetries*, *JHEP* **02** (2019) 133, [`1811.09637`].

[123] L. Di Luzio, M. Fedele, M. Giannotti, F. Mescia and E. Nardi, *Stellar evolution confronts axion models*, *JCAP* **02** (2022) 035, [`2109.10368`].

[124] L. Di Luzio, M. Giannotti, E. Nardi and L. Visinelli, *The landscape of QCD axion models*, *Phys. Rept.* **870** (2020) 1–117, [`2003.01100`].

[125] J. E. Kim, *A composite invisible axion*, *Phys. Rev. D* **31** (1985) 1733.

[126] D. B. Kaplan, *Opening the Axion Window*, *Nucl. Phys. B* **260** (1985) 215–226.

[127] K. Choi and J. E. Kim, *Dynamical Axion*, *Phys. Rev. D* **32** (1985) 1828.

[128] K. Choi, H. Kim and S. Yun, *Natural inflation with multiple sub-Planckian axions*, *Phys. Rev. D* **90** (2014) 023545, [`1404.6209`].

[129] K. Choi and S. H. Im, *Realizing the relaxion from multiple axions and its UV completion with high scale supersymmetry*, *JHEP* **01** (2016) 149, [`1511.00132`].

[130] D. E. Kaplan and R. Rattazzi, *Large field excursions and approximate discrete symmetries from a clockwork axion*, *Phys. Rev. D* **93** (2016) 085007, [`1511.01827`].

[131] A. Ernst, A. Ringwald and C. Tamarit, *Axion Predictions in $SO(10) \times U(1)_{PQ}$ Models*, *JHEP* **02** (2018) 103, [`1801.04906`].

[132] L. Di Luzio, A. Ringwald and C. Tamarit, *Axion mass prediction from minimal grand unification*, *Phys. Rev. D* **98** (2018) 095011, [`1807.09769`].

[133] A. Ernst, L. Di Luzio, A. Ringwald and C. Tamarit, *Axion properties in GUTs*, *PoS* **CORFU2018** (2019) 054, [`1811.11860`].

[134] P. Fileviez Pérez, C. Murgui and A. D. Plascencia, *The QCD Axion and Unification*, *JHEP* **11** (2019) 093, [`1908.01772`].

[135] L. Delle Rose, G. Panico, M. Redi and A. Tesi, *Gravitational Waves from Supercool Axions*, *JHEP* **04** (2020) 025, [`1912.06139`].

[136] B. Von Harling, A. Pomarol, O. Pujolàs and F. Rompineve, *Peccei-Quinn Phase Transition at LIGO*, *JHEP* **04** (2020) 195, [`1912.07587`].

[137] D. Croon, R. Houtz and V. Sanz, *Dynamical Axions and Gravitational Waves*, *JHEP* **07** (2019) 146, [`1904.10967`].

[138] C. S. Machado, W. Ratzinger, P. Schwaller and B. A. Stefanek, *Gravitational wave probes of axionlike particles*, *Phys. Rev. D* **102** (2020) 075033, [`1912.01007`].

[139] G. Ballesteros, J. Redondo, A. Ringwald and C. Tamarit, *Several Problems in Particle Physics and Cosmology Solved in One SMASH*, *Front. Astron. Space Sci.* **6** (2019) 55, [`1904.05594`].

[140] H. Georgi, D. B. Kaplan and L. Randall, *Manifesting the Invisible Axion at Low-energies*, *Phys. Lett. B* **169** (1986) 73–78.

[141] M. Bauer, M. Neubert and A. Thamm, *Collider Probes of Axion-Like Particles*, *JHEP* **12** (2017) 044, [`1708.00443`].

[142] J. Bonilla, I. Brivio, M. B. Gavela and V. Sanz, *One-loop corrections to ALP couplings*, *JHEP* **11** (2021) 168, [`2107.11392`].

[143] W. Dekens, J. de Vries and S. Shain, *CP-violating axion interactions in effective field theory*, *JHEP* **07** (2022) 014, [`2203.11230`].

[144] P. Svrcek and E. Witten, *Axions In String Theory*, *JHEP* **06** (2006) 051, [`hep-th/0605206`].

[145] A. Ringwald, *Searching for axions and ALPs from string theory*, *J. Phys. Conf. Ser.* **485** (2014) 012013, [`1209.2299`].

[146] E. Witten, *Some Properties of O(32) Superstrings*, *Phys. Lett. B* **149** (1984) 351–356.

[147] J. P. Conlon, *The QCD axion and moduli stabilisation*, *JHEP* **05** (2006) 078, [`hep-th/0602233`].

[148] M. Cicoli, M. Goodsell and A. Ringwald, *The type IIB string axiverse and its low-energy phenomenology*, *JHEP* **10** (2012) 146, [`1206.0819`].

[149] A. Arvanitaki, S. Dimopoulos, S. Dubovsky, N. Kaloper and J. March-Russell, *String Axiverse*, *Phys. Rev. D* **81** (2010) 123530, [`0905.4720`].

[150] T. W. Donnelly, S. J. Freedman, R. S. Lytel, R. D. Peccei and M. Schwartz, *Do Axions Exist?*, *Phys. Rev. D* **18** (1978) 1607.

[151] P. Sikivie, *Experimental Tests of the Invisible Axion*, *Phys. Rev. Lett.* **51** (1983) 1415–1417.

[152] C. A. J. O'Hare, "cajohare/axionlimits: Axionlimits." `https://cajohare.github.io/AxionLimits/`, July, 2020. 10.5281/zenodo.3932430.

[153] Y. Fukuda, T. Kohmoto, S. i. Nakajima and M. Kunitomo, *Production and detection of axions by using optical resonators*, *Prog. Cryst. Growth Charact. Mater.* **33** (1996) 363–366.

[154] F. Hoogeveen and T. Ziegenhagen, *Production and detection of light bosons using optical resonators*, *Nucl. Phys. B* **358** (1991) 3–26.

[155] K. Ehret et al., *New ALPS Results on Hidden-Sector Lightweights*, *Phys. Lett. B* **689** (2010) 149–155, [`1004.1313`].

[156] R. Bähre et al., *Any light particle search II —Technical Design Report*, *JINST* **8** (2013) T09001, [`1302.5647`].

[157] C. Albrecht, S. Barbanotti, H. Hintz, K. Jensch, R. Klos, W. Maschmann et al., *Straightening of Superconducting HERA Dipoles for the Any-Light-Particle-Search Experiment ALPS II*, *EPJ Tech. Instrum.* **8** (2021) 5, [`2004.13441`].

[158] A. Hallal, G. Messineo, M. D. Ortiz, J. Gleason, H. Hollis, D. B. Tanner et al., *The heterodyne sensing system for the ALPS II search for sub-eV weakly interacting particles*, *Phys. Dark Univ.* **35** (2022) 100914, [`2010.02334`].

[159] J. Jaeckel, M. Jankowiak and M. Spannowsky, *LHC probes the hidden sector*, *Phys. Dark Univ.* **2** (2013) 111–117, [`1212.3620`].

[160] K. Mimasu and V. Sanz, *ALPs at Colliders*, *JHEP* **06** (2015) 173, [`1409.4792`].

[161] J. Jaeckel and M. Spannowsky, *Probing MeV to 90 GeV axion-like particles with LEP and LHC*, *Phys. Lett. B* **753** (2016) 482–487, [`1509.00476`].

[162] I. Brivio, M. B. Gavela, L. Merlo, K. Mimasu, J. M. No, R. del Rey et al., *ALPs Effective Field Theory and Collider Signatures*, *Eur. Phys. J. C* **77** (2017) 572, [`1701.05379`].

[163] S. Knapen, T. Lin, H. K. Lou and T. Melia, *Searching for Axionlike Particles with Ultraperipheral Heavy-Ion Collisions*, *Phys. Rev. Lett.* **118** (2017) 171801, [`1607.06083`].

[164] CMS collaboration, A. M. Sirunyan et al., *Evidence for light-by-light scattering and searches for axion-like particles in ultraperipheral PbPb collisions at* $\sqrt{s_{\rm NN}} = 5.02$ *TeV*, *Phys. Lett. B* **797** (2019) 134826, [`1810.04602`].

[165] ATLAS collaboration, G. Aad et al., *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* **11** (2021) 050, [`2008.05355`].

[166] A. Mariotti, D. Redigolo, F. Sala and K. Tobioka, *New LHC bound on low-mass diphoton resonances*, *Phys. Lett. B* **783** (2018) 13–18, [`1710.01743`].

[167] M. Freytsis, Z. Ligeti and J. Thaler, *Constraining the Axion Portal with* $B \to K l^+ l^-$, *Phys. Rev. D* **81** (2010) 034001, [`0911.5355`].

[168] E. Izaguirre, T. Lin and B. Shuve, *Searching for Axionlike Particles in Flavor-Changing Neutral Current Processes*, *Phys. Rev. Lett.* **118** (2017) 111802, [`1611.09355`].

[169] 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, [`1709.00009`].

[170] X. Cid Vidal, A. Mariotti, D. Redigolo, F. Sala and K. Tobioka, *New Axion Searches at Flavor Factories*, *JHEP* **01** (2019) 113, [`1810.09452`].

[171] 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, [`1901.02031`].

[172] L. Merlo, F. Pobbe, S. Rigolin and O. Sumensari, *Revisiting the production of ALPs at B-factories*, *JHEP* **06** (2019) 091, [`1905.03259`].

[173] BELLE-II collaboration, F. Abudinén et al., *Search for Axion-Like Particles produced in e^+e^- collisions at Belle II*, *Phys. Rev. Lett.* **125** (2020) 161806, [`2007.13071`].

[174] D. Aristizabal Sierra, V. De Romeri, L. J. Flores and D. K. Papoulias, *Axionlike particles searches in reactor experiments*, *JHEP* **03** (2021) 294, [`2010.15712`].

[175] 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, [`1412.5174`].

[176] 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) 537–544, [`1810.11336`].

[177] S. Gori, G. Perez and K. Tobioka, *KOTO vs. NA62 Dark Scalar Searches*, *JHEP* **08** (2020) 110, [`2005.05170`].

[178] 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, [`2002.04623`].

[179] CHARM collaboration, F. Bergsma et al., *Search for Axion Like Particle Production in 400-GeV Proton - Copper Interactions*, *Phys. Lett. B* **157** (1985) 458–462.

[180] E. M. Riordan et al., *A Search for Short Lived Axions in an Electron Beam Dump Experiment*, *Phys. Rev. Lett.* **59** (1987) 755.

[181] J. D. Bjorken, S. Ecklund, W. R. Nelson, A. Abashian, C. Church, B. Lu et al., *Search for Neutral Metastable Penetrating Particles Produced in the SLAC Beam Dump*, *Phys. Rev. D* **38** (1988) 3375.

[182] B. Döbrich, J. Jaeckel, F. Kahlhoefer, A. Ringwald and K. Schmidt-Hoberg, *ALPtraum: ALP production in proton beam dump experiments*, *JHEP* **02** (2016) 018, [`1512.03069`].

[183] 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, [`1504.04855`].

[184] B. Döbrich, J. Jaeckel and T. Spadaro, *Light in the beam dump - ALP production from decay photons in proton beam-dumps*, *JHEP* **05** (2019) 213, [`1904.02091`].

[185] L. Darmé, F. Giacchino, E. Nardi and M. Raggi, *Invisible decays of axion-like particles: constraints and prospects*, *JHEP* **06** (2021) 009, [`2012.07894`].

[186] K. J. Kelly, S. Kumar and Z. Liu, *Heavy axion opportunities at the DUNE near detector*, *Phys. Rev. D* **103** (2021) 095002, [`2011.05995`].

[187] V. Brdar, B. Dutta, W. Jang, D. Kim, I. M. Shoemaker, Z. Tabrizi et al., *Axionlike Particles at Future Neutrino Experiments: Closing the Cosmological Triangle*, *Phys. Rev. Lett.* **126** (2021) 201801, [`2011.07054`].

[188] 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, [`1806.02348`].

[189] G. Aielli et al., *Expression of interest for the CODEX-b detector*, *Eur. Phys. J. C* **80** (2020) 1177, [`1911.00481`].

[190] NA62 collaboration, E. Cortina Gil et al., *The Beam and detector of the NA62 experiment at CERN*, *JINST* **12** (2017) P05025, [`1703.08501`].

[191] KOTO collaboration, J. K. Ahn et al., *Search for the $K_L\to\pi^0\nu\overline{\nu}$ and $K_L\to\pi^0X^0$ decays at the J-PARC KOTO experiment*, *Phys. Rev. Lett.* **122** (2019) 021802, [`1810.09655`].

[192] KLEVER Project collaboration, M. Moulson, *KLEVER: An Experiment to Measure BR($K_L \to \pi^0\nu\bar{\nu}$) at the CERN SPS*, *J. Phys. Conf. Ser.* **1526** (2020) 012028, [`1912.10037`].

[193] IAXO collaboration, E. Armengaud et al., *Physics potential of the International Axion Observatory (IAXO)*, *JCAP* **06** (2019) 047, [`1904.09155`].

[194] IAXO collaboration, A. Abeln et al., *Conceptual design of BabyIAXO, the intermediate stage towards the International Axion Observatory*, *JHEP* **05** (2021) 137, [`2010.12076`].

[195] T. Dafni, C. A. O'Hare, B. Lakić, J. Galán, F. J. Iguaz, I. G. Irastorza et al., *Weighing the solar axion*, *Phys. Rev. D* **99** (2019) 035037, [`1811.09290`].

[196] J. Jaeckel and L. J. Thormaehlen, *Distinguishing Axion Models with IAXO*, *JCAP* **1903** (2019) 039, [`1811.09278`].

[197] G. Raffelt and L. Stodolsky, *Mixing of the Photon with Low Mass Particles*, *Phys. Rev. D* **37** (1988) 1237.

[198] K. van Bibber, P. M. McIntyre, D. E. Morris and G. G. Raffelt, *A Practical Laboratory Detector for Solar Axions*, *Phys. Rev. D* **39** (1989) 2089.

[199] CAST collaboration, E. Arik et al., *Probing eV-scale axions with CAST*, *JCAP* **02** (2009) 008, [`0810.4482`].

[200] I. G. Irastorza et al., *Towards a new generation axion helioscope*, *JCAP* **1106** (2011) 013, [`1103.5334`].

[201] CAST collaboration, K. Zioutas et al., *First results from the CERN Axion Solar Telescope (CAST)*, *Phys. Rev. Lett.* **94** (2005) 121301, [`hep-ex/0411033`].

[202] CAST collaboration, S. Andriamonje et al., *An Improved limit on the axion-photon coupling from the CAST experiment*, *JCAP* **04** (2007) 010, [`hep-ex/0702006`].

[203] CAST collaboration, S. Andriamonje et al., *Search for 14.4-keV solar axions emitted in the M1-transition of Fe-57 nuclei with CAST*, *JCAP* **12** (2009) 002, [`0906.4488`].

[204] CAST collaboration, S. Andriamonje et al., *Search for solar axion emission from ^7Li and $D(p,\gamma)^3He$ nuclear decays with the CAST γ-ray calorimeter*, *JCAP* **03** (2010) 032, [`0904.2103`].

[205] CAST collaboration, S. Aune et al., *CAST search for sub-eV mass solar axions with 3He buffer gas*, *Phys. Rev. Lett.* **107** (2011) 261302, [`1106.3919`].

[206] CAST collaboration, M. Arik et al., *Search for Solar Axions by the CERN Axion Solar Telescope with ^{3}He Buffer Gas: Closing the Hot Dark Matter Gap*, *Phys. Rev. Lett.* **112** (2014) 091302, [`1307.1985`].

[207] CAST collaboration, M. Arik et al., *New solar axion search using the CERN Axion Solar Telescope with ^{4}He filling*, *Phys. Rev. D* **92** (2015) 021101, [`1503.00610`].

[208] CAST collaboration, V. Anastassopoulos et al., *New CAST Limit on the Axion-Photon Interaction*, *Nature Phys.* **13** (2017) 584–590, [`1705.02290`].

[209] IAXO collaboration, I. Irastorza et al., *The International Axion Observatory IAXO. Letter of Intent to the CERN SPS committee*, CERN-SPSC-2013-022.

[210] E. Armengaud et al., *Conceptual Design of the International Axion Observatory (IAXO)*, *JINST* **9** (2014) T05002, [`1401.3233`].

[211] IAXO collaboration, E. Armengaud et al., *Physics potential of the International Axion Observatory (IAXO)*, *JCAP* **1906** (2019) 047, [`1904.09155`].

[212] XENON collaboration, E. Aprile et al., *Excess electronic recoil events in XENON1T*, *Phys. Rev. D* **102** (2020) 072004, [`2006.09721`].

[213] L. Di Luzio, M. Fedele, M. Giannotti, F. Mescia and E. Nardi, *Solar axions cannot explain the XENON1T excess*, *Phys. Rev. Lett.* **125** (2020) 131804, [`2006.12487`].

[214] C. Gao, J. Liu, L.-T. Wang, X.-P. Wang, W. Xue and Y.-M. Zhong, *Reexamining the Solar Axion Explanation for the XENON1T Excess*, *Phys. Rev. Lett.* **125** (2020) 131806, [`2006.14598`].

[215] E. Aprile et al., *Search for New Physics in Electronic Recoil Data from XENONnT*, [`2207.11330`].

[216] ADMX collaboration, N. Du et al., *A Search for Invisible Axion Dark Matter with the Axion Dark Matter Experiment*, *Phys. Rev. Lett.* **120** (2018) 151301, [`1804.05750`].

[217] C. B. Adams et al., *Axion Dark Matter*, in *2022 Snowmass Summer Study*, March, 2022, [`2203.14923`].

[218] HAYSTAC collaboration, K. M. Backes et al., *A quantum-enhanced search for dark matter axions*, *Nature* **590** (2021) 238–242, [`2008.01853`].

[219] A. V. Dixit, S. Chakram, K. He, A. Agrawal, R. K. Naik, D. I. Schuster et al., *Searching for Dark Matter with a Superconducting Qubit*, *Phys. Rev. Lett.* **126** (2021) 141302, [`2008.12231`].

[220] MADMAX collaboration, P. Brun et al., *A new experimental approach to probe QCD axion dark matter in the mass range above 40 μeV*, *Eur. Phys. J. C* **79** (2019) 186, [`1901.07401`].

[221] M. Lawson, A. J. Millar, M. Pancaldi, E. Vitagliano and F. Wilczek, *Tunable axion plasma haloscopes*, *Phys. Rev. Lett.* **123** (2019) 141802, [`1904.11872`].

[222] BREAD collaboration, J. Liu et al., *Broadband Solenoidal Haloscope for Terahertz Axion Detection*, *Phys. Rev. Lett.* **128** (2022) 131801, [`2111.12103`].

[223] 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, [`1212.2970`].

[224] Y. Kahn, B. R. Safdi and J. Thaler, *Broadband and Resonant Approaches to Axion Dark Matter Detection*, *Phys. Rev. Lett.* **117** (2016) 141801, [`1602.01086`].

[225] J. L. Ouellet et al., *First Results from ABRACADABRA-10 cm: A Search for Sub-μeV Axion Dark Matter*, *Phys. Rev. Lett.* **122** (2019) 121802, [`1810.12257`].

[226] 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, [`1306.6089`].

[227] D. Aybas et al., *Search for Axionlike Dark Matter Using Solid-State Nuclear Magnetic Resonance*, *Phys. Rev. Lett.* **126** (2021) 141802, [`2101.01241`].

[228] F. Takahashi, M. Yamada and W. Yin, *XENON1T Excess from Anomaly-Free Axionlike Dark Matter and Its Implications for Stellar Cooling Anomaly*, *Phys. Rev. Lett.* **125** (2020) 161801, [`2006.10035`].

[229] D. A. Dicus, E. W. Kolb, V. L. Teplitz and R. V. Wagoner, *Astrophysical Bounds on the Masses of Axions and Higgs Particles*, *Phys. Rev. D* **18** (1978) 1829.

[230] J. R. Ellis and K. A. Olive, *Constraints on Light Particles From Stellar Evolution*, *Nucl. Phys. B* **223** (1983) 252–268.

[231] G. G. Raffelt, *Astrophysical axion bounds diminished by screening effects*, *Phys. Rev. D* **33** (1986) 897.

[232] G. G. Raffelt and D. S. P. Dearborn, *Bounds on Weakly Interacting Particles From Observational Lifetimes of Helium Burning Stars*, *Phys. Rev. D* **37** (1988) 549–551.

[233] G. G. Raffelt, *Stars as laboratories for fundamental physics: The astrophysics of neutrinos, axions, and other weakly interacting particles.* University of Chicago Press, May, 1996.

[234] H. An, M. Pospelov and J. Pradler, *New stellar constraints on dark photons*, *Phys. Lett. B* **725** (2013) 190–195, [`1302.3884`].

[235] M. Giannotti, *Hints of new physics from stars*, *PoS* **ICHEP2016** (2016) 076, [`1611.04651`].

[236] M. Giannotti, I. G. Irastorza, J. Redondo, A. Ringwald and K. Saikawa, *Stellar Recipes for Axion Hunters*, *JCAP* **1710** (2017) 010, [`1708.02111`].

[237] G. Lucente, O. Straniero, P. Carenza, M. Giannotti and A. Mirizzi, *Constraining Heavy Axionlike Particles by Energy Deposition in Globular Cluster Stars*, *Phys. Rev. Lett.* **129** (2022) 011101, [`2203.01336`].

[238] J. Isern, M. Hernanz and E. Garcia-Berro, *Axion cooling of white dwarfs*, *Astrophys. J.* **392** (1992) L23.

[239] J. Isern, E. Garcia-Berro, S. Torres and S. Catalan, *Axions and the cooling of white dwarf stars*, *Astrophys. J.* **682** (2008) L109, [`0806.2807`].

[240] J. Isern, S. Catalan, E. Garcia-Berro and S. Torres, *Axions and the white dwarf luminosity function*, *J. Phys. Conf. Ser.* **172** (2009) 012005, [`0812.3043`].

[241] M. M. Miller Bertolami, *Limits on the neutrino magnetic dipole moment from the luminosity function of hot white dwarfs*, *Astron. Astrophys.* **562** (2014) A123, [`1407.1404`].

[242] 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* **1410** (2014) 069, [`1406.7712`].

[243] A. Bischoff-Kim, M. H. Montgomery and D. E. Winget, *Strong limits on the DFSZ axion mass with G117-B15A*, *Astrophys. J.* **675** (2008) 1512, [`0711.2041`].

[244] A. H. Córsico, L. G. Althaus, M. M. Miller Bertolami, A. D. Romero, E. Garcia-Berro, J. Isern et al., *The rate of cooling of the pulsating white dwarf star G117−B15A: a new asteroseismological inference of the axion mass*, *Mon. Not. Roy. Astron. Soc.* **424** (2012) 2792, [`1205.6180`].

[245] A. H. Córsico, L. G. Althaus, A. D. Romero, A. S. Mukadam, E. Garcia-Berro, J. Isern et al., *An independent limit on the axion mass from the variable white dwarf star R548*, *JCAP* **1212** (2012) 010, [`1211.3389`].

[246] A. H. Córsico, A. D. Romero, L. G. Althaus, E. García-Berro, J. Isern, S. O. Kepler et al., *An asteroseismic constraint on the mass of the axion from the period drift of the pulsating DA white dwarf star L19-2*, *JCAP* **1607** (2016) 036, [`1605.06458`].

[247] T. Battich, A. H. Córsico, L. G. Althaus and M. M. Miller Bertolami, *First axion bounds from a pulsating helium-rich white dwarf star*, *JCAP* **1608** (2016) 062, [`1605.07668`].

[248] C. Dessert, A. J. Long and B. R. Safdi, *X-ray Signatures of Axion Conversion in Magnetic White Dwarf Stars*, *Phys. Rev. Lett.* **123** (2019) 061104, [`1903.05088`].

[249] C. Dessert, D. Dunsky and B. R. Safdi, *Upper limit on the axion-photon coupling from magnetic white dwarf polarization*, *Phys. Rev. D* **105** (2022) 103034, [`2203.04319`].

[250] G. G. Raffelt and D. S. P. Dearborn, *Bounds on Hadronic Axions From Stellar Evolution*, *Phys. Rev. D* **36** (1987) 2211.

[251] 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, [`1406.6053`].

[252] O. Straniero, A. Ayala, M. Giannotti, A. Mirizzi and I. Domínguez, *Axion-Photon Coupling: Astrophysical Constraints*, in *Proceedings, 11th Patras Workshop on Axions, WIMPs and WISPs (Axion-WIMP 2015): Zaragoza, Spain, June 22-26, 2015*, pp. 77–81, 2015, DOI.

[253] N. Viaux, M. Catelan, P. B. Stetson, G. Raffelt, J. Redondo, A. A. R. Valcarce et al., *Neutrino and axion bounds from the globular cluster M5 (NGC 5904)*, *Phys. Rev. Lett.* **111** (2013) 231301, [`1311.1669`].

[254] R. C. Dohm-Palmer and E. D. Skillman, *The ratio of blue to red supergiants in sextans a from hst imaging*, [`astro-ph/0203284`].

[255] K. B. W. McQuinn, E. D. Skillman, J. J. Dalcanton, A. E. Dolphin, J. Holtzman, D. R. Weisz et al., *Observational Constraints on Red and Blue Helium Burning Sequences*, *Astrophys. J.* **740** (2011) 48, [`1108.1405`].

[256] A. Friedland, M. Giannotti and M. Wise, *Constraining the Axion-Photon Coupling with Massive Stars*, *Phys. Rev. Lett.* **110** (2013) 061101, [`1210.1271`].

[257] G. Carosi, A. Friedland, M. Giannotti, M. J. Pivovaroff, J. Ruz and J. K. Vogel, *Probing the axion-photon coupling: phenomenological and experimental perspectives. A snowmass white paper*, in *Proceedings, 2013 Community Summer Study on the Future of U.S. Particle Physics: Snowmass on the Mississippi (CSS2013): Minneapolis, MN, USA, July 29-August 6, 2013*, 2013, [`1309.7035`].

[258] C. O. Heinke and W. C. G. Ho, *Direct Observation of the Cooling of the Cassiopeia A Neutron Star*, *Astrophys. J.* **719** (2010) L167–L171, [`1007.4719`].

[259] L. B. Leinson, *Axion mass limit from observations of the neutron star in Cassiopeia A*, *JCAP* **1408** (2014) 031, [`1405.6873`].

[260] KAMIOKANDE-II collaboration, K. Hirata et al., *Observation of a Neutrino Burst from the Supernova SN 1987a*, *Phys. Rev. Lett.* **58** (1987) 1490–1493.

[261] R. M. Bionta et al., *Observation of a Neutrino Burst in Coincidence with Supernova SN 1987a in the Large Magellanic Cloud*, *Phys. Rev. Lett.* **58** (1987) 1494.

[262] M. Aglietta et al., *On the event observed in the Mont Blanc Underground Neutrino observatory during the occurrence of Supernova 1987a*, *Europhys. Lett.* **3** (1987) 1315–1320.

[263] M. S. Turner, *Axions from SN 1987a*, *Phys. Rev. Lett.* **60** (1988) 1797.

[264] G. G. Raffelt, *Astrophysical axion bounds*, *Lect. Notes Phys.* **741** (2008) 51–71, [`hep-ph/0611350`].

[265] P. Carenza, 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, [`1906.11844`].

[266] T. J. Loredo and D. Q. Lamb, *Bayesian analysis of neutrinos observed from supernova SN-1987A*, *Phys. Rev. D* **65** (2002) 063002, [`astro-ph/0107260`].

[267] M. Giannotti, L. D. Duffy and R. Nita, *New constraints for heavy axion-like particles from supernovae*, *JCAP* **1101** (2011) 015, [`1009.5714`].

[268] F. Calore, P. Carenza, M. Giannotti, J. Jaeckel and A. Mirizzi, *Bounds on axionlike particles from the diffuse supernova flux*, *Phys. Rev. D* **102** (2020) 123005, [`2008.11741`].

[269] A. De Angelis, M. Roncadelli and O. Mansutti, *Evidence for a new light spin-zero boson from cosmological gamma-ray propagation?*, *Phys. Rev. D* **76** (2007) 121301, [`0707.4312`].

[270] A. Mirizzi, G. G. Raffelt and P. D. Serpico, *Signatures of axion-like particles in the spectra of TeV gamma-ray sources*, *Phys. Rev. D* **76** (2007) 023001, [`0704.3044`].

[271] M. Simet, D. Hooper and P. D. Serpico, *The Milky Way as a Kiloparsec-Scale Axionscope*, *Phys. Rev. D* **77** (2008) 063001, [`0712.2825`].

[272] A. De Angelis, O. Mansutti, M. Persic and M. Roncadelli, *Photon propagation and the VHE gamma-ray spectra of blazars: how transparent is really the Universe?*, *Mon. Not. Roy. Astron. Soc.* **394** (2009) L21–L25, [`0807.4246`].

[273] M. A. Sanchez-Conde, D. Paneque, E. Bloom, F. Prada and A. Dominguez, *Hints of the existence of Axion-Like-Particles from the gamma-ray spectra of cosmological sources*, *Phys. Rev. D* **79** (2009) 123511, [`0905.3270`].

[274] A. Dominguez, M. A. Sanchez-Conde and F. Prada, *Axion-like particle imprint in cosmological very-high-energy sources*, *JCAP* **1111** (2011) 020, [`1106.1860`].

[275] A. De Angelis, G. Galanti and M. Roncadelli, *Relevance of axion-like particles for very-high-energy astrophysics*, *Phys. Rev. D* **84** (2011) 105030, [`1106.1132`].

[276] W. Essey and A. Kusenko, *On weak redshift dependence of gamma-ray spectra of distant blazars*, *Astrophys. J.* **751** (2012) L11, [`1111.0815`].

[277] D. Horns and M. Meyer, *Indications for a pair-production anomaly from the propagation of VHE gamma-rays*, *JCAP* **1202** (2012) 033, [`1201.4711`].

[278] M. Meyer, D. Horns and M. Raue, *First lower limits on the photon-axion-like particle coupling from very high energy gamma-ray observations*, *Phys. Rev. D* **87** (2013) 035027, [`1302.1208`].

[279] G. I. Rubtsov and S. V. Troitsky, *Breaks in gamma-ray spectra of distant blazars and transparency of the Universe*, *JETP Lett.* **100** (2014) 355–359, [`1406.0239`].

[280] S. V. Troitsky, *Axion-like particles and the propagation of gamma rays over astronomical distances*, *JETP Lett.* **105** (2017) 55–59, [`1612.01864`].

[281] K. Kohri and H. Kodama, *Axion-Like Particles and Recent Observations of the Cosmic Infrared Background Radiation*, *Phys. Rev. D* **96** (2017) 051701, [`1704.05189`].

[282] A. Korochkin, G. Rubtsov and S. Troitsky, *Search for anomalous features in gamma-ray blazar spectra corrected for the absorption on the extragalactic background light*, *JCAP* **12** (2019) 002, [`1810.03443`].

[283] Y.-F. Liang, C. Zhang, Z.-Q. Xia, L. Feng, Q. Yuan and Y.-Z. Fan, *Constraints on axion-like particle properties with TeV gamma-ray observations of Galactic sources*, *JCAP* **06** (2019) 042, [`1804.07186`].

[284] D. A. Sanchez, S. Fegan and B. Giebels, *Evidence for a cosmological effect in γ-ray spectra of BL Lacs*, *Astron. Astrophys.* **554** (2013) A75, [`1303.5923`].

[285] A. Domínguez and M. Ajello, *Spectral analysis of Fermi-LAT blazars above 50 GeV*, *Astrophys. J.* **813** (2015) L34, [`1510.07913`].

[286] J. Biteau and D. A. Williams, *The extragalactic background light, the Hubble constant, and anomalies: conclusions from 20 years of TeV gamma-ray observations*, *Astrophys. J.* **812** (2015) 60, [`1502.04166`].

[287] P. Agrawal, K. V. Berghaus, J. Fan, A. Hook, G. Marques-Tavares and T. Rudelius, *Some open questions in axion theory*, in *2022 Snowmass Summer Study*, March, 2022, [`2203.08026`].

[288] P. Agrawal, J. Fan, M. Reece and L.-T. Wang, *Experimental Targets for Photon Couplings of the QCD Axion*, *JHEP* **02** (2018) 006, [`1709.06085`].

[289] K. Fraser and M. Reece, *Axion Periodicity and Coupling Quantization in the Presence of Mixing*, *JHEP* **05** (2020) 066, [`1910.11349`].

[290] M. A. Buen-Abad, J. Fan, M. Reece and C. Sun, *Challenges for an axion explanation of the muon $g-2$ measurement*, *JHEP* **09** (2021) 101, [`2104.03267`].

[291] R. S. Gupta, Z. Komargodski, G. Perez and L. Ubaldi, *Is the Relaxion an Axion?*, *JHEP* **02** (2016) 166, [`1509.00047`].

[292] A. Ringwald and K. Saikawa, *Axion dark matter in the post-inflationary Peccei-Quinn symmetry breaking scenario*, *Phys. Rev. D* **93** (2016) 085031, [`1512.06436`].

[293] R. Foot, H. Lew and R. R. Volkas, *Electric charge quantization*, *J. Phys. G* **19** (1993) 361–372, [`hep-ph/9209259`].

[294] R. A. Bertlmann, *Anomalies in Quantum Field Theory.* Oxford University Press, Nov., 2000, 10.1093/acprof:oso/9780198507628.001.0001.

[295] MoEDAL collaboration, B. Acharya et al., *Search for magnetic monopoles with the MoEDAL prototype trapping detector in 8 TeV proton-proton collisions at the LHC*, *JHEP* **08** (2016) 067, [`1604.06645`].

[296] M. Chala, G. Guedes, M. Ramos and J. Santiago, *Running in the ALPs*, *Eur. Phys. J. C* **81** (2021) 181, [`2012.09017`].

[297] M. Srednicki, *Axion Couplings to Matter. 1. CP Conserving Parts*, *Nucl. Phys. B* **260** (1985) 689–700.

[298] D. B. Kaplan and A. Manohar, *Strange Matrix Elements in the Proton from Neutral Current Experiments*, *Nucl. Phys. B* **310** (1988) 527–547.

[299] C. Grojean, E. E. Jenkins, A. V. Manohar and M. Trott, *Renormalization Group Scaling of Higgs Operators and \Gamma(h -> \gamma \gamma)*, *JHEP* **04** (2013) 016, [`1301.2588`].

[300] J. Ellis, *TikZ-Feynman: Feynman diagrams with TikZ*, *Comput. Phys. Commun.* **210** (2017) 103–123, [`1601.05437`].

[301] H. H. Patel, *Package-X: A Mathematica package for the analytic calculation of one-loop integrals*, *Comput. Phys. Commun.* **197** (2015) 276–290, [`1503.01469`].

[302] A. Latosinski, K. A. Meissner and H. Nicolai, *Axions without Peccei-Quinn Symmetry*, [`1010.5417`].

[303] A. Latosinski, K. A. Meissner and H. Nicolai, *Neutrino Loops from Neutrino Mixing*, [`1112.0134`].

[304] A. Latosinski, K. A. Meissner and H. Nicolai, *Neutrino Mixing and the Axion-Gluon Vertex*, *Nucl. Phys. B* **868** (2013) 596–626, [`1203.3886`].

[305] R. Z. Ferreira, M. C. D. Marsh and E. Müller, *Strong supernovae bounds on ALPs from quantum loops*, [`2205.07896`].

[306] A. Hees, O. Minazzoli, E. Savalle, Y. V. Stadnik and P. Wolf, *Violation of the equivalence principle from light scalar dark matter*, *Phys. Rev. D* **98** (2018) 064051, [`1807.04512`].

[307] P. W. Graham, D. E. Kaplan and S. Rajendran, *Cosmological Relaxation of the Electroweak Scale*, *Phys. Rev. Lett.* **115** (2015) 221801, [`1504.07551`].

[308] M. Bauer, G. Rostagni and J. Spinner, *The Axion-Higgs Portal*, [`2207.05762`].

[309] G. Buchalla, A. J. Buras and M. E. Lautenbacher, *Weak decays beyond leading logarithms*, *Rev. Mod. Phys.* **68** (1996) 1125–1144, [`hep-ph/9512380`].

[310] 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, [`1704.06639`].

[311] 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, [`1709.04486`].

[312] W. Buchmuller and D. Wyler, *Effective Lagrangian Analysis of New Interactions and Flavor Conservation*, *Nucl. Phys. B* **268** (1986) 621–653.

[313] B. Grzadkowski, M. Iskrzynski, M. Misiak and J. Rosiek, *Dimension-Six Terms in the Standard Model Lagrangian*, *JHEP* **10** (2010) 085, [`1008.4884`].

[314] F. Feruglio, *The Chiral approach to the electroweak interactions*, *Int. J. Mod. Phys. A* **8** (1993) 4937–4972, [`hep-ph/9301281`].

[315] 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) 330–335, [`1212.3305`].

[316] G. Buchalla, O. Catà and C. Krause, *Complete Electroweak Chiral Lagrangian with a Light Higgs at NLO*, *Nucl. Phys. B* **880** (2014) 552–573, [`1307.5017`].

[317] K. Choi, S. H. Im, C. B. Park and S. Yun, *Minimal Flavor Violation with Axion-like Particles*, *JHEP* **11** (2017) 070, [`1708.00021`].

[318] L. Calibbi, F. Goertz, D. Redigolo, R. Ziegler and J. Zupan, *Minimal axion model from flavor*, *Phys. Rev. D* **95** (2017) 095009, [`1612.08040`].

[319] C. D. Froggatt and H. B. Nielsen, *Hierarchy of Quark Masses, Cabibbo Angles and CP Violation*, *Nucl. Phys. B* **147** (1979) 277–298.

[320] C. Cornella, P. Paradisi and O. Sumensari, *Hunting for ALPs with Lepton Flavor Violation*, *JHEP* **01** (2020) 158, [`1911.06279`].

[321] P. Escribano and A. Vicente, *Ultralight scalars in leptonic observables*, *JHEP* **03** (2021) 240, [`2008.01099`].

[322] E949, E787 collaboration, S. Adler et al., *Measurement of the K+ –> pi+ nu nu branching ratio*, *Phys. Rev. D* **77** (2008) 052003, [`0709.1000`].

[323] NA62 collaboration, R. Fantechi, *The NA62 experiment at CERN: status and perspectives*, in *12th Conference on Flavor Physics and CP Violation*, July, 2014, [`1407.8213`].

[324] NA62 collaboration, E. Cortina Gil et al., *Search for a feebly interacting particle X in the decay $K^+ \to \pi^+ X$*, *JHEP* **03** (2021) 058, [`2011.11329`].

[325] STAR collaboration, B. I. Abelev et al., *Azimuthal Charged-Particle Correlations and Possible Local Strong Parity Violation*, *Phys. Rev. Lett.* **103** (2009) 251601, [`0909.1739`].

[326] STAR collaboration, B. I. Abelev et al., *Observation of charge-dependent azimuthal correlations and possible local strong parity violation in heavy ion collisions*, *Phys. Rev. C* **81** (2010) 054908, [`0909.1717`].

[327] B. Batell, M. Pospelov and A. Ritz, *Multi-lepton Signatures of a Hidden Sector in Rare B Decays*, *Phys. Rev. D* **83** (2011) 054005, [`0911.4938`].

[328] M. Bauer, M. Neubert, S. Renner, M. Schnubel and A. Thamm, *The Low-Energy Effective Theory of Axions and ALPs*, *JHEP* **04** (2021) 063, [`2012.12272`].

[329] L. J. Hall and M. B. Wise, *Flavor changing Higgs boson couplings*, *Nucl. Phys. B* **187** (1981) 397–408.

[330] S. Oda, Y. Shoji and D.-S. Takahashi, *High Scale Validity of the DFSZ Axion Model with Precision*, *JHEP* **03** (2020) 011, [`1912.01147`].

[331] F. Kling, S. Su and W. Su, *2HDM Neutral Scalars under the LHC*, *JHEP* **06** (2020) 163, [`2004.04172`].

[332] C. A. J. O'Hare and E. Vitagliano, *Cornering the axion with CP-violating interactions*, *Phys. Rev. D* **102** (2020) 115026, [`2010.03889`].

[333] L. Di Luzio, R. Gröber and P. Paradisi, *Hunting for CP-violating axionlike particle interactions*, *Phys. Rev. D* **104** (2021) 095027, [`2010.13760`].

[334] J. Hisano, K. Tsumura and M. J. S. Yang, *QCD Corrections to Neutron Electric Dipole Moment from Dimension-six Four-Quark Operators*, *Phys. Lett. B* **713** (2012) 473–480, [`1205.2212`].

[335] CAST collaboration, V. Anastassopoulos et al., *New CAST Limit on the Axion-Photon Interaction*, *Nature Phys.* **13** (2017) 584–590, [`1705.02290`].

[336] ADMX collaboration, N. Du et al., *A Search for Invisible Axion Dark Matter with the Axion Dark Matter Experiment*, *Phys. Rev. Lett.* **120** (2018) 151301, [`1804.05750`].

[337] ADMX collaboration, T. Braine et al., *Extended Search for the Invisible Axion with the Axion Dark Matter Experiment*, *Phys. Rev. Lett.* **124** (2020) 101303, [`1910.08638`].

[338] T. Vonk, F.-K. Guo and U.-G. Meißner, *Precision calculation of the axion-nucleon coupling in chiral perturbation theory*, *JHEP* **03** (2020) 138, [`2001.05327`].

[339] 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, [`1803.00993`].

[340] F. Ertas and F. Kahlhoefer, *On the interplay between astrophysical and laboratory probes of MeV-scale axion-like particles*, *JHEP* **07** (2020) 050, [`2004.01193`].

[341] M. Giannotti, I. Irastorza, J. Redondo and A. Ringwald, *Cool WISPs for stellar cooling excesses*, *JCAP* **05** (2016) 057, [`1512.08108`].

[342] 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, [`1810.07192`].

[343] N. Bar, K. Blum and G. D'Amico, *Is there a supernova bound on axions?*, *Phys. Rev. D* **101** (2020) 123025, [`1907.05020`].

[344] P. Cigan et al., *High angular resolution ALMA images of dust and molecules in the SN 1987A ejecta*, *Astrophys. J.* **886** (2019) 51, [`1910.02960`].

[345] 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, [`2004.06078`].

[346] J. Jaeckel and L. J. Thormaehlen, *Axions as a probe of solar metals*, *Phys. Rev. D* **100** (2019) 123020, [`1908.10878`].

[347] C. A. J. O'Hare, A. Caputo, A. J. Millar and E. Vitagliano, *Axion helioscopes as solar magnetometers*, *Phys. Rev. D* **102** (2020) 043019, [`2006.10415`].

[348] D. M. Lazarus, G. C. Smith, R. Cameron, A. C. Melissinos, G. Ruoso, Y. K. Semertzidis et al., *A Search for solar axions*, *Phys. Rev. Lett.* **69** (1992) 2333–2336.

[349] S. Moriyama, M. Minowa, T. Namba, Y. Inoue, Y. Takasu and A. Yamamoto, *Direct search for solar axions by using strong magnetic field and x-ray detectors*, *Phys. Lett. B* **434** (1998) 147, [`hep-ex/9805026`].

[350] Y. Inoue, T. Namba, S. Moriyama, M. Minowa, Y. Takasu, T. Horiuchi et al., *Search for sub-electronvolt solar axions using coherent conversion of axions into photons in magnetic field and gas helium*, *Phys. Lett. B* **536** (2002) 18–23, [`astro-ph/0204388`].

[351] Y. Inoue, Y. Akimoto, R. Ohta, T. Mizumoto, A. Yamamoto and M. Minowa, *Search for solar axions with mass around 1 eV using coherent conversion of axions into photons*, *Phys. Lett. B* **668** (2008) 93–97, [`0806.2230`].

[352] K. Barth et al., *CAST constraints on the axion-electron coupling*, *JCAP* **1305** (2013) 010, [`1302.6283`].

[353] G. G. Raffelt, *Plasmon Decay Into Low Mass Bosons in Stars*, *Phys. Rev. D* **37** (1988) 1356.

[354] J. Jaeckel, E. Masso, J. Redondo, A. Ringwald and F. Takahashi, *The Need for purely laboratory-based axion-like particle searches*, *Phys. Rev. D* **75** (2007) 013004, [`hep-ph/0610203`].

[355] S. A. Bludman and K. A. van Riper, *Equation of state of an ideal Fermi gas.*, *Astrophys. J.* **212** (March, 1977) 859–872.

[356] A. Caputo, A. J. Millar and E. Vitagliano, *Revisiting longitudinal plasmon-axion conversion in external magnetic fields*, *Phys. Rev. D* **101** (2020) 123004, [`2005.00078`].

[357] E. Guarini, P. Carenza, J. Galan, M. Giannotti and A. Mirizzi, *Production of axionlike particles from photon conversions in large-scale solar magnetic fields*, *Phys. Rev. D* **102** (2020) 123024, [`2010.06601`].

[358] L. D. Landau, *On the vibrations of the electronic plasma*, *J. Phys. (USSR)* **10** (1946) 25–34.

[359] M. Krief, A. Feigel and D. Gazit, *Solar opacity calculations using the super-transition-array method*, *Astrophys. J.* **821** (2016) 45, [`1601.01930`].

[360] S. Couvidat, S. Turck-Chieze and A. G. Kosovichev, *Solar seismic models and the neutrino predictions*, *Astrophys. J.* **599** (2003) 1434–1448, [`astro-ph/0203107`].

[361] E. Vitagliano, J. Redondo and G. Raffelt, *Solar neutrino flux at keV energies*, *JCAP* **12** (2017) 010, [`1708.02248`].

[362] P. Carenza and G. Lucente, *Revisiting axion-electron bremsstrahlung emission rates in astrophysical environments*, *Phys. Rev. D* **103** (2021) 123024, [`2104.09524`].

[363] A. N. Cox and J. N. Stewart, *Radiative and Conductive Opacities for Eleven Astrophysical Mixtures.*, *Astrophys. J. Suppl.* **11** (Jan., 1965) 22.

[364] F. J. Rogers and C. A. Iglesias, *Radiative atomic Rosseland mean opacity tables*, *Astrophys. J. Suppl.* **79** (1992) 507–568.

[365] N. R. Badnell, M. A. Bautista, K. Butler, F. Delahaye, C. Mendoza, P. Palmeri et al., *Up-dated opacities from the Opacity Project*, *Mon. Not. Roy. Astron. Soc.* **360** (2005) 458–464, [`astro-ph/0410744`].

[366] N. H. Magee, J. Abdallah, J., R. E. H. Clark, J. S. Cohen, L. A. Collins, G. Csanak et al., *Atomic Structure Calculations and New LOS Alamos Astrophysical Opacities*, in *Astrophysical Applications of Powerful New Databases* (S. J. Adelman and W. L. Wiese, eds.), vol. 78 of *Astronomical Society of the Pacific Conference Series*, p. 51, Jan., 1995.

[367] C. Blancard, P. Cossé and G. Faussurier, *Solar Mixture Opacity Calculations Using Detailed Configuration and Level Accounting Treatments*, *Astrophys. J.* **745** (Jan., 2012) 10.

[368] Borexino collaboration, G. Bellini et al., *Search for Solar Axions Produced in $p(d,{}^3\mathrm{He})A$ Reaction with Borexino Detector*, *Phys. Rev. D* **85** (2012) 092003, [`1203.6258`].

[369] A. Bhusal, N. Houston and T. Li, *Searching for Solar Axions Using Data from the Sudbury Neutrino Observatory*, *Phys. Rev. Lett.* **126** (2021) 091601, [`2004.02733`].

[370] S. Moriyama, *A Proposal to search for a monochromatic component of solar axions using Fe-57*, *Phys. Rev. Lett.* **75** (1995) 3222–3225, [`hep-ph/9504318`].

[371] M. Krcmar, Z. Krecak, M. Stipcevic, A. Ljubicic and D. A. Bradley, *Search for invisible axions using Fe-57*, *Phys. Lett. B* **442** (1998) 38, [`nucl-ex/9801005`].

[372] Y. M. Gavrilyuk et al., *First result of the experimental search for the 9.4 keV solar axion reactions with ${}^{83}Kr$ in the copper proportional counter*, *Phys. Part. Nucl.* **46** (2015) 152–156, [`1405.1271`].

[373] R. Röhlsberger, *Nuclear Condensed Matter Physics with Synchrotron Radiation.* Springer-Verlag Berlin Heidelberg, Germany, 2004.

[374] M. Asplund, N. Grevesse, A. J. Sauval and P. Scott, *The chemical composition of the Sun*, *Ann. Rev. Astron. Astrophys.* **47** (2009) 481–522, [`0909.0948`].

[375] W. C. Haxton and K. Y. Lee, *Red giant evolution, metallicity and new bounds on hadronic axions*, *Phys. Rev. Lett.* **66** (1991) 2557–2560.

[376] F. T. Avignone, C. Baktash, W. C. Barker, F. P. Calaprice, R. W. Dunford, W. C. Haxton et al., *Search for Axions From the 1115-kev Transition of ${}^{65}Cu$*, *Phys. Rev. D* **37** (1988) 618–630.

[377] F. T. Avignone, R. J. Creswick, J. D. Vergados, P. Pirinen, P. C. Srivastava and J. Suhonen, *Estimating the flux of the 14.4 keV solar axions*, *JCAP* **01** (2018) 021, [`1711.06979`].

[378] CUORE collaboration, F. Alessandria et al., *Search for 14.4 keV solar axions from M1 transition of Fe-57 with CUORE crystals*, *JCAP* **05** (2013) 007, [`1209.2800`].

[379] N. Vinyoles, A. M. Serenelli, F. L. Villante, S. Basu, J. Bergström, M. C. Gonzalez-Garcia et al., *A new Generation of Standard Solar Models*, *Astrophys. J.* **835** (2017) 202, [`1611.09867`].

[380] M. J. Seaton, *OP data on CD for mean opacities and radiative accelerations*, *Mon. Not. Roy. Astron. Soc.* **362** (2005) 1, [`astro-ph/0411010`].

[381] D. R. Williams, "Earth fact sheet." `https://nssdc.gsfc.nasa.gov/planetary/factsheet/earthfact.html`, 2020.

[382] J. N. Bahcall, S. Basu, M. Pinsonneault and A. M. Serenelli, *Helioseismological implications of recent solar abundance determinations*, *Astrophys. J.* **618** (2005) 1049–1056, [`astro-ph/0407060`].

[383] H. M. Antia and S. Basu, *The Discrepancy between solar abundances and helioseismology*, *Astrophys. J.* **620** (2005) L129–L132, [`astro-ph/0501129`].

[384] C. Pena-Garay and A. Serenelli, *Solar neutrinos and the solar composition problem*, [`0811.2424`].

[385] A. Serenelli, S. Basu, J. W. Ferguson and M. Asplund, *New Solar Composition: The Problem With Solar Models Revisited*, *Astrophys. J.* **705** (2009) L123–L127, [`0909.2668`].

[386] C. Mendoza, *Computation of Atomic Astrophysical Opacities*, *Atoms* **6** (2018) 28, [`1704.03528`].

[387] F. L. Villante and A. Serenelli, *The relevance of nuclear reactions for Standard Solar Models construction*, *Front. Astron. Space Sci.* **7** (2021) 112, [`2101.03077`].

[388] A. M. Serenelli, *New Results on Standard Solar Models*, *Astrophys. Space Sci.* **328** (2010) 13–21, [`0910.3690`].

[389] J. N. Bahcall, S. Basu and M. H. Pinsonneault, *How uncertain are solar neutrino predictions?*, *Phys. Lett. B* **433** (1998) 1–8, [`astro-ph/9805135`].

[390] J. N. Bahcall, M. H. Pinsonneault and S. Basu, *Solar models: Current epoch and time dependences, neutrinos, and helioseismological properties*, *Astrophys. J.* **555** (2001) 990–1012, [`astro-ph/0010346`].

[391] J. N. Bahcall and M. H. Pinsonneault, *What do we (not) know theoretically about solar neutrino fluxes?*, *Phys. Rev. Lett.* **92** (2004) 121301, [`astro-ph/0402114`].

[392] J. N. Bahcall, A. M. Serenelli and S. Basu, *New solar opacities, abundances, helioseismology, and neutrino fluxes*, *Astrophys. J. Lett.* **621** (2005) L85–L88, [`astro-ph/0412440`].

[393] J. Colgan, D. P. Kilcrease, N. H. Magee, M. E. Sherrill, J. Abdallah, J., P. Hakel et al., *A New Generation of Los Alamos Opacity Tables*, *Astrophys. J.* **817** (Feb., 2016) 116, [`1601.01005`].

[394] G. Mondet, C. Blancard, P. Cossé and G. Faussurier, *Opacity Calculations for Solar Mixtures*, *Astrophys. J. Suppl.* **220** (Sep., 2015) 2.

[395] J. N. Bahcall, A. M. Serenelli and S. Basu, *10,000 standard solar models: a Monte Carlo simulation*, *Astrophys. J. Suppl.* **165** (2006) 400–431, [`astro-ph/0511337`].

[396] A. Weiss and H. Schlattl, *GARSTEC—the Garching Stellar Evolution Code*, *Astrophysics and Space Science* **316** (2008) 99–106.

[397] A. Kramida, Yu. Ralchenko, J. Reader and and NIST ASD Team. NIST Atomic Spectra Database (ver. 5.8), [Online]. Available: `https://physics.nist.gov/asd` [2021, January 19]. National Institute of Standards and Technology, Gaithersburg, MD., 2020.

[398] T. Abe, K. Hamaguchi and N. Nagata, *Atomic Form Factors and Inverse Primakoff Scattering of Axion*, *Phys. Lett. B* **815** (2021) 136174, [`2012.02508`].

[399] J. H. Hubbell, W. J. Veigele, E. A. Briggs, R. T. Brown, D. T. Cromer and R. J. Howerton, *Atomic form factors, incoherent scattering functions, and photon scattering cross sections*, *J. Phys. Chem. Ref. Data* **4** (1975) 471–538.

[400] G. Elwert, *Verschärfte Berechnung von Intensität und Polarisation im kontinuierlichen Röntgenspektrum*, *Annalen der Physik* **426** (Jan., 1939) 178–208.

[401] N. Vinyoles, A. Serenelli, F. L. Villante, S. Basu, J. Redondo and J. Isern, *New axion and hidden photon constraints from a solar data global fit*, *JCAP* **10** (2015) 015, [`1501.01639`].

[402] G. Raffelt and D. Seckel, *Bounds on Exotic Particle Interactions from SN 1987a*, *Phys. Rev. Lett.* **60** (1988) 1793.

[403] G. G. Raffelt, *Astrophysical methods to constrain axions and other novel particle phenomena*, *Phys. Rept.* **198** (1990) 1–113.

[404] M. S. Turner, *Dirac neutrinos and SN1987A*, *Phys. Rev. D* **45** (1992) 1066–1075.

[405] G. Raffelt and D. Seckel, *A selfconsistent approach to neutral current processes in supernova cores*, *Phys. Rev. D* **52** (1995) 1780–1799, [`astro-ph/9312019`].

[406] W. Keil, H.-T. Janka, D. N. Schramm, G. Sigl, M. S. Turner and J. R. Ellis, *A Fresh look at axions and SN-1987A*, *Phys. Rev. D* **56** (1997) 2419–2432, [`astro-ph/9612222`].

[407] T. Fischer, S. Chakraborty, M. Giannotti, A. Mirizzi, A. Payez and A. Ringwald, *Probing axions with the neutrino signal from the next galactic supernova*, *Phys. Rev. D* **94** (2016) 085012, [`1605.08780`].

[408] T. Fischer, P. Carenza, B. Fore, M. Giannotti, A. Mirizzi and S. Reddy, *Observable signatures of enhanced axion emission from protoneutron stars*, *Phys. Rev. D* **104** (2021) 103012, [`2108.13726`].

[409] J. Keller and A. Sedrakian, *Axions from cooling compact stars*, *Nucl. Phys. A* **897** (2013) 62–69, [`1205.6940`].

[410] A. Sedrakian, *Axion cooling of neutron stars*, *Phys. Rev. D* **93** (2016) 065044, [`1512.07828`].

[411] K. Hamaguchi, N. Nagata, K. Yanagi and J. Zheng, *Limit on the Axion Decay Constant from the Cooling Neutron Star in Cassiopeia A*, *Phys. Rev. D* **98** (2018) 103015, [`1806.07151`].

[412] M. V. Beznogov, E. Rrapaj, D. Page and S. Reddy, *Constraints on Axion-like Particles and Nucleon Pairing in Dense Matter from the Hot Neutron Star in HESS J1731-347*, *Phys. Rev. C* **98** (2018) 035802, [`1806.07991`].

[413] A. Sedrakian, *Axion cooling of neutron stars. II. Beyond hadronic axions*, *Phys. Rev. D* **99** (2019) 043011, [`1810.00190`].

[414] L. B. Leinson, *Impact of axions on the Cassiopea A neutron star cooling*, *JCAP* **09** (2021) 001, [`2105.14745`].

[415] A. Abeln et al., *Axion search with BabyIAXO in view of IAXO*, *PoS* **ICHEP2020** (2021) 631, [`2012.06634`].

[416] S. Aune et al., *Low background x-ray detection with Micromegas for axion research*, *JINST* **9** (2014) P01001, [`1310.3391`].

[417] S. Andriamonje et al., *Development and performance of Microbulk Micromegas detectors*, *JINST* **5** (2010) P02001.

[418] J. Castel, S. Cebrián, T. Dafni, J. Galán, I. G. Irastorza, G. Luzón et al., *Status of low mass WIMP detector TREX-DM*, *J. Phys. Conf. Ser.* **1312** (2019) 012010.

[419] P. Lechner, C. Fiorini, A. Longoni, G. Lutz, A. Pahlke, H. Soltau et al., *Silicon drift detectors for high resolution, high count rate x-ray spectroscopy at room temperature*, *Powder Diffraction* **47** (June, 2003) 53–58.

[420] F. E. Christensen, A. C. Jakobsen, N. F. Brejnholt, K. K. Madsen, A. Hornstrup, N. J. Westergaard et al., *Coatings for the NuSTAR mission*, in *Optics for EUV, X-Ray, and Gamma-Ray Astronomy V* (S. L. O'Dell and G. Pareschi, eds.), vol. 8147, pp. 298 – 316, International Society for Optics and Photonics, SPIE, 2011, DOI.

[421] NuSTAR collaboration, F. A. Harrison et al., *The Nuclear Spectroscopic Telescope Array (NuSTAR) High-Energy X-Ray Mission*, *Astrophys. J.* **770** (2013) 103, [`1301.7307`].

[422] F. A. Harrison, W. W. Craig, F. E. Christensen, C. J. Hailey, W. W. Zhang, S. E. Boggs et al., *The Nuclear Spectroscopic Telescope Array (NuSTAR) High-energy X-Ray Mission*, *apj* **770** (June, 2013) 103, [`1301.7307`].

[423] C. M. H. Chen, W. R. Cook, F. A. Harrison, J. Y. Y. Lin, P. H. Mao and S. M. Schindler, *Characterization of the HEFT CdZnTe pixel detectors*, in *Hard X-Ray and Gamma-Ray Detector Physics V* (L. A. Franks, A. Burger, R. B. James and P. L. Hink, eds.), vol. 5198, pp. 9 – 18, International Society for Optics and Photonics, SPIE, 2004, DOI.

[424] M. Freitas, F. Medeiros and E. M. Yoshimura, *Detection properties of cdznte semiconductor for diagnostic x-ray spectroscopic applications*, in *Cross-Disciplinary Applied Research in Materials Science and Technology*, vol. 480 of *Materials Science Forum*, pp. 53–58, Trans Tech Publications Ltd, Jan., 2005, DOI.

[425] D. Unger, A. Abeln, C. Enss, A. Fleischmann, D. Hengstler, S. Kempf et al., *High-resolution for IAXO: MMC-based x-ray detectors*, *J. Instrum.* **16** (June, 2021) P06006.

[426] F. S. Porter, J. Gygax, R. L. Kelley, C. A. Kilbourne, J. M. King, P. Beiersdorfer et al., *Performance of the ebit calorimeter spectrometer*, *Review of Scientific Instruments* **79** (Oct., 2008) .

[427] L. Darmé, L. Di Luzio, M. Giannotti and E. Nardi, *Selective enhancement of the QCD axion couplings*, *Phys. Rev. D* **103** (2021) 015034, [`2010.15846`].

[428] C. Bobeth, T. Ewerth, F. Kruger and J. Urban, *Analysis of neutral Higgs boson contributions to the decays* $\bar{B}(\ s^{)} \to \ell^+\ell^-$ *and* $\bar{B} \to K\ell^+\ell^-$, *Phys. Rev. D* **64** (2001) 074014, [`hep-ph/0104284`].

[429] H. E. Logan and U. Nierste, $B_{s,d} \to \ell^+\ell^-$ *in a two Higgs doublet model*, *Nucl. Phys. B* **586** (2000) 39–55, [`hep-ph/0004139`].

[430] L. D. Landau and E. M. Lifshitz, *Statistical Physics.* Pergamon Press, Oxford, 2 ed., 1968.

Danksagung

Mein großer Dank gilt meinem Doktorvater Prof. Jörg Jäckel, der mir in den vergangenen Jahren stets mit hilfreichem Rat zur Seite stand. Seine Begeisterung für die Physik und seine Art, sicher Geglaubtes in Frage zu stellen, haben mein wissenschaftliches Denken und diese Arbeit nachhaltig beeinflusst.

Zudem danke ich Prof. Jan Pawlowski für die Bereitschaft das Zweitgutachten für diese Doktorarbeit zu verfassen.

Meinen Co-Autoren danke ich für die produktive Zusammenarbeit, den anregenden Austausch und die Möglichkeit auch in der Pandemie Wissenschaft als kooperative Arbeit zu erfahren.

Sehr dankbar bin ich der gesamten Phänomenologie Arbeitsgruppe am Institut für theoretische Physik für eine angenehme Arbeitsatmosphäre, kollegiale Zusammenarbeit und gegenseitige Motivation. Ein besonderer Dank gilt meinen langjährigen Bürokollegen Gonzalo, Alaric und Paul für sowohl hilfreiche als auch unterhaltsame Diskussionen.

Dem Graduiertenkolleg *Particle Physics beyond the Standard Model* der Deutschen Forschungsgemeinschaft danke ich für die finanzielle und inhaltliche Unterstützung dieser Arbeit und für den engen Austausch zwischen Experimentalphysikern und Theoretikern, von dem ich sehr profitiert habe.

Zu besonderem Dank verpflichtet bin ich Ruben, Alaric und Valentina für aufmerksames Korrekturlesen dieser Arbeit und hilfreiche Kommentare.

Prof. Klaus Hasselmann und Susanne Hasselmann möchte ich für die hilfreichen Gespräche vor der Studienwahl und die Vermittlung an das Deutsche Zentrum für Luft- und Raumfahrt danken. Dies hat mir die ersten Einblicke in aktive Forschung ermöglicht.

Danken möchte ich auch meinen Physik- und Mathematiklehrern, insbesondere Herrn Wendt, Frau Schwertfeger und Herrn Wohlenberg, die bereits früh mein Interesse wecken konnten und aufzeigten, dass es in der Physik stets noch viel Spannendes zu entdecken gibt.

Meinen Freunden und Studienkollegen danke ich für neun wundervolle Jahre in Heidelberg, für Ablenkung und Zerstreuung, wenn immer es vonnöten war und dafür, meinen Studienort zu einem Zuhause gemacht zu haben.

Ganz besonders dankbar bin ich meinen Eltern und meiner Schwester für ihre unnachgiebige Unterstützung, für unsere ausführlichen Gespräche und unseren starken und währenden Zusammenhalt. Ihnen sei diese Arbeit gewidmet.