RECEIVED: October 25, 2018

REVISED: February 14, 2019

ACCEPTED: March 11, 2019

PUBLISHED: March 28, 2019

Axion global fits with Peccei-Quinn symmetry breaking before inflation using GAMBIT

Sebastian Hoof,^a Felix Kahlhoefer,^b Pat Scott,^a Christoph Weniger^c
and Martin White^d

^a*Department of Physics, Imperial College London, Blackett Laboratory,
Prince Consort Road, London SW7 2AZ, U.K.*

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

^c*GRAPPA, Institute of Physics, University of Amsterdam,
Science Park 904, 1098 XH Amsterdam, Netherlands*

^d*Department of Physics, University of Adelaide,
Adelaide, SA 5005, Australia*

E-mail: s.hoof15@imperial.ac.uk, kahlhoefer@physik.rwth-aachen.de,
p.scott@imperial.ac.uk, c.weniger@uva.nl,
martin.white@adelaide.edu.au

ABSTRACT: We present global fits of cosmologically stable axion-like particle and QCD axion models in the mass range 0.1 neV to 10 eV. We focus on the case where the Peccei-Quinn symmetry is broken before the end of inflation, such that the initial value of the axion field can be considered to be homogeneous throughout the visible Universe. We include detailed likelihood functions from light-shining-through-wall experiments, haloscopes, helioscopes, the axion relic density, horizontal branch stars, supernova 1987A, white dwarf cooling, and gamma-ray observations. We carry out both frequentist and Bayesian analyses, with and without the inclusion of white dwarf cooling. We explore the degree of fine-tuning present in different models and identify parameter regions where it is possible for QCD axion models to account for both the dark matter in the Universe and the cooling hints, comparing them to specific DFSZ- and KSVZ-type models. We find the most credible parameter regions, allowing us to set (prior-dependent) upper and lower bounds on the axion mass. Our analysis also suggests that QCD axions in this scenario most probably make up a non-negligible but sub-dominant component of the dark matter in the Universe.

KEYWORDS: Beyond Standard Model, Cosmology of Theories beyond the SM

ARXIV EPRINT: [1810.07192](https://arxiv.org/abs/1810.07192)

Contents

1	Introduction	1
2	Axion physics	2
2.1	The QCD axion	2
2.2	Axion-like particles	4
2.3	Axion creation mechanisms	5
2.3.1	Realignment mechanism	5
2.3.2	Topological defects	7
2.3.3	Thermal creation	7
3	Axion models and theoretical uncertainties	7
3.1	General ALP model	9
3.2	QCD axion models	9
3.2.1	KSVZ models	12
3.2.2	DFSZ models	13
3.3	ALP models with constant mass	13
4	Observables, experiments and likelihoods	13
4.1	Light-shining-through-wall experiments	13
4.2	Helioscopes	16
4.2.1	2007 CAST results	17
4.2.2	2017 CAST results	18
4.3	Haloscopes (cavity experiments)	19
4.3.1	RBF results	20
4.3.2	UF results	21
4.3.3	ADMX results 1998–2009	21
4.3.4	ADMX results 2018	22
4.4	Dark matter relic density	22
4.5	Astrophysical probes	23
4.5.1	Distortions of gamma-ray spectra	23
4.5.2	Supernova 1987A	25
4.5.3	Horizontal branch stars and R parameter	27
4.5.4	White dwarf cooling hints	27
5	Results and discussion	30
5.1	Sampling algorithms and settings	30
5.2	General ALP models	30
5.3	QCD axions	33
5.4	DFSZ- and KSVZ-type models	39
5.5	Cooling hints	42

5.5.1	QCD axions	43
5.5.2	DFSZ- and KSVZ-type axions	47
6	Conclusions	50
A	Integrating solar models for the signal prediction in CAST	52
B	Numerical implementation of the solution to the axion field equation	54
C	H.E.S.S. likelihood implementation details	56
D	Prior-dependence of results	57
E	Overview of new capabilities	60

1 Introduction

QCD axions [1–4] and axion-like particles (ALPs) are perhaps among the most intriguing classes of hypothetical particles. They have been extensively studied in the literature and experiments are expected to probe the relevant parameter space for many axion models in the near future (see ref. [5] for an up-to-date review on axion searches). Because QCD axions can behave as cold dark matter (DM) [6–9], the particularly interesting regions of the parameter space are where they contribute significantly to the DM density of the Universe. If so, they would solve at least two outstanding problems in physics at the same time (the other being the Strong CP problem [1, 2]).

There are now many complementary searches for axions¹ underway, and many new search strategies have emerged in recent years (see ref. [5]). Axions could also explain some apparent anomalies in the cooling of white dwarf stars [10–16], or the transparency of the Universe to gamma rays [17–22]. It is therefore crucial to combine all available results in order to extract the maximum information from the data. In doing so, we can learn more about the parameter space of different models, help guide the planning of future searches towards the most promising search areas and — if axions do indeed exist — find them in the correlated signals of several experiments and determine their properties.

Some of these goals can be satisfactorily achieved by over-plotting exclusion limits from several experiments. However, such simple exclusion plots generally have some shortcomings. They make assumptions about the relative importance of interactions with other matter, and about important inputs such as solar physics, the local density of DM and theoretical uncertainties. Incompatibilities between the assumptions can undermine the validity of a simple combination. They also do not allow for a quantitative model comparison between different theories that include axions.

¹In this paper, we use the generic term “axion” to refer to both QCD axions and ALPs. We use more specific terms such as “QCD axion” to make more model-specific statements.

A solution is to perform a consistent global statistical analysis of all available constraints, accounting for the leading theoretical and experimental uncertainties. Implementing such an analysis requires careful design, along with compromises relating to the availability of data and the total computational runtime required. An example of such a computational framework for BSM physics is **GAMBIT** [23]. Full details of **GAMBIT**'s features can be found in the relevant publications [23–28] and physics analyses [29–34]. The most relevant features for this work are **GAMBIT**'s modularity, and the options that it offers for carrying out both Bayesian and frequentist analyses. **GAMBIT**'s structure allows easy integration of new components such as models, theory calculations, likelihoods and sampling algorithms. **GAMBIT** contains a variety of advanced samplers for both Bayesian and frequentist analyses, which are particularly useful for including nuisance parameters and assessing fine-tuning.

In the following section, we review some aspects of axion physics, including specific issues that should be accounted for in global fits. In section 3, we turn to axion models, their corresponding effective field theories and the family tree of axion models available in **GAMBIT**. Section 4 describes our observables and likelihood functions, including their implementation, incorporated experimental data, potential caveats and restrictions. Results and discussion of the first global scans of axion models can be found in section 5, and section 6 summarises our results.

The axion routines developed for this paper are available in the DarkBit [25] module of **GAMBIT** 1.3.0, available at <https://gambit.hepforge.org> under the 3-clause BSD license.² Likelihood and posterior samples from this study can be downloaded from Zenodo [35].

2 Axion physics

Here we present a brief overview of axion physics, highlighting caveats of our current implementation and pointing out opportunities for future extensions. More details can be found in the literature [5, 36–40].

2.1 The QCD axion

The Peccei-Quinn (PQ) mechanism [1, 2] is a proposed solution to the Strong CP problem that gives rise to a pseudo-scalar particle: the QCD axion [3, 4]. The QCD axion is an excellent DM candidate [6–9], as it can account for the entire cosmological abundance of DM via the *vacuum misalignment* or *realignment* mechanism (section 2.3.1). The original QCD axion model also inspired further archetypical models, such as the KSVZ [41, 42] and DFSZ [43, 44] axion models, which we will introduce in Secs 3.2.1 and 3.2.2.

The symmetries of the Standard Model (SM) Lagrangian permit a term of the form

$$\mathcal{L}_{\text{QCD}} \supset -\frac{\alpha_S}{8\pi} \theta_{\text{QCD}} G_{\mu\nu}^a \tilde{G}^{\mu\nu,a}, \quad (2.1)$$

where $G_{\mu\nu}^a$ is the gluon field strength, a is the SU(3) gauge index and α_S is the strong coupling constant. The angle $\theta_{\text{QCD}} \in [-\pi, \pi]$ is an unknown parameter and $\tilde{G}^{\mu\nu,a} =$

²<http://opensource.org/licenses/BSD-3-Clause>

$\epsilon^{\mu\nu\kappa\lambda} G_{\kappa\lambda}^a/2$. A contribution is also generated by chiral transformations due to the chiral anomaly, which replaces θ_{QCD} by an effective angle,

$$\theta_{\text{eff}} \equiv \theta_{\text{QCD}} - \arg [\det(Y_d Y_u)] , \quad (2.2)$$

where Y_d and Y_u are the down- and up-type Yukawa matrices, respectively [45, section 29.5]. The $G\tilde{G}$ -term is anti-symmetric under the discrete parity (P) and charge-parity (CP) transformations. Due to the presence of this term, one would naïvely expect the strong interaction to show some CP -violating effects — especially because weak interactions are known to violate P maximally and CP mildly. There is no CP -violation if and only if θ_{eff} vanishes (modulo the periodicity).

Experiments trying to measure the CP -violating electric dipole moment of the neutron (nEDM) can place limits on the value of θ_{eff} . Within 40–50% uncertainty, the dipole moment induced by the $G\tilde{G}$ -term is given by [46, 47]

$$|d_n| \approx (2.4 \times 10^{-16} \text{ e cm}) |\theta_{\text{eff}}| . \quad (2.3)$$

The current limit on the nEDM is $|d_n| < 3.6 \times 10^{-26} \text{ e cm}$ at 95% confidence level [48], resulting in $|\theta_{\text{eff}}| \lesssim 10^{-10}$.

This observation poses a fine-tuning issue in the SM, commonly referred to as the Strong CP problem. The mechanism suggested by Peccei and Quinn [1, 2] solves this problem by introducing a new global, axial $U(1)$ symmetry. This so-called PQ symmetry is spontaneously broken by the vacuum expectation value v of a complex scalar field. The resulting Nambu-Goldstone boson is the QCD axion, a pseudo-scalar field denoted by $a(x)$. It adds another contribution to θ_{eff} of the form $Na(x)/v$, where the non-zero integer N is the colour anomaly of the PQ symmetry. The associated shift symmetry can then be used to cancel the θ_{eff} term by driving $\theta_{\text{eff}} + Na(x)/v$ to zero. In fact, it can be shown by the Vafa-Witten theorem [49, 50] that the axion dynamically and asymptotically relaxes θ_{eff} to the CP -conserving minimum. The Strong CP problem is therefore effectively solved by promoting θ_{eff} to a dynamical degree of freedom.

The continuous shift symmetry of the PQ $U(1)$ phase forbids any mass terms for the axion at the Lagrangian level. In the presence of an $a G\tilde{G}$ term in the Lagrangian, however, this symmetry is broken after the QCD phase transition due to topologically non-trivial fluctuations of the gluon fields, leaving only a discrete shift symmetry of size $2\pi Nv$.

The resulting effective potential can be written as³

$$V(a) = f_a^2 m_a^2 [1 - \cos(a/f_a)] , \quad (2.4)$$

where m_a is the temperature-dependent axion mass and $f_a \equiv v/N$. The axion is therefore practically massless until the time of the QCD phase transition, where it picks up a small mass that can be calculated using the chiral Lagrangian formalism.

³The full potential [51, 52] includes corrections to the simple cosine shape. There is evidence that the potential in (2.4) is a good approximation at higher temperatures [53, 54], relevant for solving the axion field equation in section 2.3.1.

The zero-temperature mass of the QCD axion was initially calculated by relating v to the scale of weak interactions [3]. However, this model was ruled out and the parameter space for f_a was opened up in order to make the axion “invisible” — in the sense that it evaded experimental constraints at the time. Neglecting suppressed quark mass ratios, the QCD axion mass at zero temperature is given by [3]

$$m_{a,0} \simeq \frac{\sqrt{z_d}}{1+z_d} \frac{f_{\pi^0}}{f_a} m_{\pi^0}, \quad (2.5)$$

where $z_d = m_u/m_d$ is the up- and down-quark mass ratio, and f_{π^0} and m_{π^0} are the decay constant and mass of the pion respectively. Using next-to-leading order chiral perturbation theory, a recent study found the QCD axion mass at zero temperature to be [52]

$$m_{a,0} = 5.70(7) \text{ } \mu\text{eV} \left(\frac{10^{12} \text{ GeV}}{f_a} \right). \quad (2.6)$$

At temperatures exceeding the QCD scale, the axion becomes increasingly light as the shift symmetry is restored. Numerical estimates of the temperature dependence can be obtained directly from recent lattice QCD simulations [54, 55]. These results are well-approximated at higher temperatures by a power law $m_a(T) \propto T^{-\beta/2}$ for some $\beta > 0$, and by a constant axion mass below some transition temperature T_χ . This also agrees with the behaviour predicted from the finite-temperature dilute instanton gas approximation [56].

As we solve all relevant equations numerically in this paper, it would be straightforward to include the full lattice QCD results and calculate the axion mass at every given temperature. However, including the statistical uncertainties⁴ is most easily achieved by having a parametrised form of the temperature dependence of the QCD axion mass. This is commonly done by introducing two parameters, β and T_χ , which we fit to lattice QCD results. More details about the implementation can be found in section 3.2.

2.2 Axion-like particles

Apart from the original QCD axion, pseudo-Nambu-Goldstone bosons with fundamental shift symmetries appear in many other contexts. These are usually referred to as ALPs (axion-like particles), and are appealing for theorists and model builders because of the shift symmetry, and the ability of their field values to undergo relaxation via the realignment mechanism in the same manner as axions. They do not necessarily solve the Strong CP problem. However, under the right circumstances, they can be cold DM candidates, or play an important role in solving other physics puzzles. ALPs typically arise from the breaking of a U(1) symmetry at some scale f_a and generate a mass from explicit breaking of the residual symmetry at scale Λ . As a result, they can be fairly light, with masses of the order $m_a \sim \Lambda^2/f_a$, and have suppressed couplings to the SM. Due to the lack of a direct relation between the two scales involved (f_a and m_a), they occupy a larger parameter space than QCD axions. More details on the theory and phenomenology of these particles can be found in the literature [5, 36, 57].

⁴From lattice QCD results, we can infer that (at zero temperature) the statistical uncertainties in the axion mass are about twice as important as the systematic ones [54].

2.3 Axion creation mechanisms

QCD axions (with roughly μeV to meV masses) and many ALP models are expected to be lighter than other hypothetical particles, such as WIMPs, which typically have GeV to TeV -scale masses [58]. For thermal DM, this would be problematic, as DM would not be sufficiently cold today to reproduce the observed large-scale structure of the Universe. However, although a small population of axions is produced thermally, the relic abundance is typically dominated by non-thermal mechanisms, namely the realignment mechanism and topological defects. In this work, we focus on the realignment mechanism, which allows axions to be both ultra-cold and very light at the same time.

2.3.1 Realignment mechanism

The equation of motion for a homogeneous QCD axion or ALP field $\theta(t) = a(t)/f_a$ with potential $V(\theta)$ in a Friedmann-Robertson-Walker-Lemaître universe reads [59, appendix B.12]

$$\ddot{\theta} + 3H(t)\dot{\theta} + \frac{1}{f_a^2} \frac{\delta V(\theta)}{\delta \theta} = 0, \quad (2.7)$$

where $H \equiv \dot{a}/a$ is the Hubble parameter. For the canonical axion potential,

$$V(\theta) = f_a^2 m_a^2 [1 - \cos(\theta)], \quad (2.8)$$

the equation of motion becomes

$$\ddot{\theta} + 3H(t)\dot{\theta} + m_a^2(t) \sin(\theta) = 0. \quad (2.9)$$

The general form of this equation does not possess analytic solutions. In the early Universe $m_a \ll H$ and the system is an overdamped oscillator. We can integrate the differential equation with boundary conditions $\theta(t_i) = \theta_i$ and $\dot{\theta}(t_i) = 0$, where the value θ_i of the axion field is called the *initial misalignment angle*.

At later times, around the time when $m_a \sim H$, the system becomes critically damped and the field starts to oscillate. In the regime of $m_a \gg H$, the axion field oscillations are adiabatic while their amplitude continuously decreases due to Hubble damping. In other words, at some time t_* , the axion field evolution is guaranteed to eventually enter an epoch where $|\theta(t)| \ll 1$ for all $t > t_*$. In this harmonic limit, (2.9) is a damped harmonic oscillator, and the field evolution is very well described by the WKB approximation [38]

$$\theta(t) \propto [m_a(t) a^3(t)]^{-\frac{1}{2}} \cos \left(\int_{t_*}^t m_a(\tau) d\tau + \delta_* \right), \quad (2.10)$$

where we can match the solution at time t_* with the phase δ_* . This limiting form of the axion scalar field oscillations can be used to demonstrate that they behave like cold DM at sufficiently late times by averaging (2.10) over an oscillation period to find the effective behaviour of the energy density, ε , such that [38]

$$\langle \varepsilon \rangle \propto m_a a^{-3}. \quad (2.11)$$

Once the mass becomes constant, this is the scaling behaviour of cold matter. We can rephrase this statement for non-relativistic axions to see that the (averaged) *comoving* axion number density is conserved:

$$n_a^{\text{com}} \equiv \frac{\rho a^3}{m_a} \simeq \frac{\varepsilon a^3}{m_a}. \quad (2.12)$$

In fact, there is an adiabatic invariant in the harmonic limit of $\sin(\theta) \simeq \theta$, which has been studied in the literature for QCD axions and ALPs [7, 9, 60]. However, if the initial misalignment angle $\theta_i \sim \mathcal{O}(1)$, then the adiabatic invariant is not sufficient to describe the system, and anharmonic effects arising from the full potential must be taken into account. These have been estimated and calculated by several authors [9, 61–64].

The energy density in axions today can hence be parametrised by an overall transfer function F , with the properties that F is bounded from below, and becomes constant in the harmonic limit $|\theta_i| \ll 1$ [9, 62, 63]. This allows us to conveniently write the axion energy density in the Universe today as

$$\rho_a = F(|\theta_i|; \beta, T_\chi, m_{a,0}) \theta_i^2, \quad (2.13)$$

where we explicitly denote the dependence of the transfer function F on the values of the axion parameters. For a given set of these parameters, the transfer function — and hence the axion relic density — could be determined once and used in subsequent calculations. For sampling over the parameter space, it is necessary to repeat this process for each set of parameters (or to tabulate the results on a sufficiently dense grid ahead of the scan).

We have thus far only considered the *homogeneous* field equation (2.9). It is not obvious that this is sufficient for determining the energy density in axions today. Let us consider the two possible scenarios in a cosmology with inflation: either PQ symmetry breaks *after* inflation, or it breaks *before* inflation ends.

In the first case, the Universe consists of a large number of causally disconnected “bubbles” where the initial misalignment angles take random values from a uniform distribution on the interval $-\pi$ to π [9]. As a consequence, the initial misalignment angle θ_i effectively becomes a function of space. Nonetheless, the resulting overall energy density in axions is fixed, because it can be calculated as the average over all the misalignment angles in all the patches. As θ_i is constant within the patches, we may use (2.13) to calculate the result:

$$\rho_a^{\text{post-inf}} = \frac{1}{2\pi} \int_{-\pi}^{\pi} F(|\theta_i|) \theta_i^2 d\theta_i = \frac{1}{\pi} \int_0^{\pi} F(|\theta_i|) \theta_i^2 d\theta_i. \quad (2.14)$$

In the second case, where the PQ symmetry breaks *before* the end of inflation, one causally connected patch gets blown up to at least the size of the observable Universe. The axion field is therefore homogeneous, with a random initial value in the interval $(-\pi, \pi]$. The stochastic nature of the initial misalignment angle gives rise to a physically motivated prior probability for θ_i in a Bayesian analysis, which will be discussed in section 5.

A major issue with this scenario is perturbations from inflation, which will affect the initial misalignment angle by some amount that depends on the energy scale of inflation [see e.g. 61]. It has been argued that this scenario is therefore finely-tuned to a degree that

can be considered “worse” than the Strong CP problem itself [65, 66]. Since the necessary additional implementation of inflationary models is beyond the scope of this work, we neglect the issue of field fluctuations during inflation.

2.3.2 Topological defects

Because axions arise from the breaking of a U(1) symmetry, topological defects known as cosmic strings and domain walls can appear [67, 68]. Decay of these defects to axions will increase the axion DM density. They have been studied extensively in the literature, but the relative importance of their contribution is still not firmly established [69–85].

Some authors have parametrised the topological defect contribution as $\Omega_a = \Omega_a^{\text{realign}}(1 + \alpha)$, where $\Omega_a^{\text{realign}}$ is the contribution from realignment. Including α as a nuisance parameter in our work would remove much of the predictability of the model, and would assume similar scaling of the energy density contributions. As we consider PQ symmetry to break before inflation, we assume that contributions from topological defects are diluted away during inflation [86].

2.3.3 Thermal creation

Axions are also thermally produced in the early Universe [68, 87]. The resulting abundance is however rather dependent on the additional new field content associated with the axion.

For the QCD axion, the Boltzmann equations for the least model-dependent processes give useful estimates. The contributions of low-temperature processes such as $\pi + \pi \rightleftharpoons \pi + a$ have been computed [68, 87–89], including all combinations of pionic states π . Thermal axion production during reheating has also been studied [90–92]. As hadronisation would not yet have taken place by the time of reheating, these calculations consider processes like axion-gluon interactions, i.e. $g + g \rightleftharpoons g + a$. Recent calculations also considered model-dependent processes with quarks q , i.e. $q + \bar{q} \rightleftharpoons g + a$ and $q + g \rightleftharpoons q + a$, which can be dominant in the early Universe [93].

A significant abundance of thermal axions modifies the effective number of relativistic species, which can be used to set limits on the axion mass [94–102]. Generally speaking, thermally-produced QCD axions with a mass of more than $\mathcal{O}(\text{eV})$ are hot DM with a relic abundance comparable to that of neutrinos and photons. Those axion models are therefore excluded, as they would exceed the observational bounds on the fraction of DM that can be hot [e.g. 40, 68, 86]. However, such bounds have some dependence on the choice of cosmological datasets [101], and the limits can be relaxed if the Universe had a non-standard thermal history [103].

3 Axion models and theoretical uncertainties

Let us now discuss the various axion models that we consider in this work with a focus on the interactions between axions and other particles. We already saw that axions acquire a mass, so that they interact gravitationally. QCD axions also interact via the strong interaction. For generic ALPs, however, it is not clear *a priori* what the PQ charges of SM

particles are, nor whether the axion interacts with a given particle at tree level or only at higher order in perturbation theory.

In order to study axion phenomenology and to identify useful observables, axions have been studied in an effective field theory (EFT) framework [104–106], which can be adjusted to a given energy scale and scenario. One can then apply exclusion limits on the effective interactions to specific axion models that establish a relation between the effective couplings and the fundamental axion parameters; examples of such models include the so-called KSVZ model [41, 42] and the DFSZ model [43, 44].

One important consequence of the fundamental shift symmetry is that axions can only directly interact with matter via derivative couplings. Below the electroweak scale, the effective Lagrangian for axion-SM interactions is

$$\mathcal{L}_a = - \sum_f \frac{g_{aff}}{2m_f} \bar{f} \gamma^\mu \gamma_5 f \partial_\mu a - \frac{g_{a\gamma\gamma}}{4} a F_{\mu\nu} \tilde{F}^{\mu\nu} - \frac{\alpha_S}{8\pi} \frac{a}{f_a} G_{\mu\nu}^b \tilde{G}^{\mu\nu,b}, \quad (3.1)$$

where a is the axion field, $\tilde{F}$ and $\tilde{G}$ are the duals of electromagnetic and strong field strengths and g_{aff} and $g_{a\gamma\gamma}$ are effective coupling constants of mass dimension -1 . In principle, f runs over all SM fermions with mass m_f , but the couplings most relevant for axion searches are those to electrons and nucleons, the latter arising from matching the EFT to chiral perturbation theory [106].⁵ Also note that the Lagrangian (3.1) is flavour-diagonal, which is not necessarily the case in all models. While we do not consider the possibility of flavour non-diagonal interactions in this paper, the resulting models can be phenomenologically interesting, rather predictive, and within reach of existing and future experimental searches [e.g. 108–110].

We now turn to the specific axion models implemented in GAMBIT (see figure 1). We describe these models in the following subsections, while we leave the discussion of the parameter ranges and prior distributions for section 5. Note that the models presented here are only a subset of the many general and more specific ALP models studied in the literature; detailed overviews can be found in review articles [5, 40].

Models in GAMBIT are defined as collections of named parameters. All relevant observables for a model must be computable from these parameters. Models can have relationships to other models, allowing parameter combinations in one model to be translated to equivalent combinations in another model or to an alternative parameterisation of the same model. This is achieved by adding new “children” to the family tree of a more general model and defining a translation between the two models. This translation may, e.g., fix the values of some of the “parent” parameters or compute them as functions of (possibly new) parameters in the child model. This ensures flexibility in the choice of the independent model parameters to work with in any given theory calculation, allowing the most convenient definitions to always be used for any calculation. More details on the general implementation of models in GAMBIT can be found in ref. [23].

⁵The same matching also gives rise to interactions between axions and mesons, which can in particular induce flavour-changing rare decays. While of phenomenological relevance for heavier ALPs [107], these processes do not lead to relevant constraints on the parameter space considered in the present work.

Figure 1. Family tree of axion models in GAMBIT. The numbers in brackets refer to the number of model parameters; $(n + m)$ indicates n (largely unconstrained) fundamental parameters of the model and m (typically well-constrained) nuisance parameters.

3.1 General ALP model

We define a new family of GAMBIT models for axions and ALPs. On top of the family tree is the GeneralALP model, with seven parameters. This model describes an effective Lagrangian that is a subset of (3.1),

$$\mathcal{L}_a^{\text{int}} = -\frac{g_{a\gamma\gamma}}{4} a F_{\mu\nu} \tilde{F}^{\mu\nu} - \frac{g_{aee}}{2m_e} \bar{e} \gamma^\mu \gamma_5 e \partial_\mu a. \quad (3.2)$$

In this study we do not include couplings of axions to nucleons. However, future versions of GAMBIT may include additional interactions such as these, subject to the availability of interesting observables and constraints.

We employ the rescaled field value $\theta = a/f_a$, assume the canonical cosine potential (2.8) for all types of axions,⁶ and take

$$m_a(T) = m_{a,0} \begin{cases} 1 & \text{if } T \leq T_\chi \\ \left(\frac{T_\chi}{T}\right)^{\beta/2} & \text{otherwise} \end{cases}. \quad (3.3)$$

A summary of all the model parameters can be found in table 1.

3.2 QCD axion models

The QCDAXion model is a child of the GeneralALP model. It is inspired by the original QCD axion models, which solve the Strong CP problem. The scale of the explicit breaking

⁶Note that this is a non-trivial assumption, as the potential could be any periodic function of θ ; more general potentials have been invoked to e.g. construct models of inflation [111]. Allowing a different shape of the potential would also require a modification of the relic density calculator, presented in appendix B.

<i>Parameter</i>	<i>Description</i>	<i>Comment</i>
$g_{a\gamma\gamma}$	Effective axion-photon coupling	Units of GeV^{-1}
g_{aee}	Effective axion-electron coupling	
f_a	Axion decay constant	Units of GeV
$m_{a,0}$	Axion zero-temperature mass	Units of eV
T_χ	Transition temperature in the broken power law for m_a	Units of MeV
β	Exponent of the broken power law for m_a	$\beta > 0$
θ_i	Initial misalignment angle	$-\pi < \theta_i \leq \pi$
Λ	Breaking scale of the residual symmetry	Units of MeV
C_{aee}	Axion-electron coupling	
Λ_χ	Zero-temperature topological susceptibility, $\Lambda_\chi^4 \equiv \chi(T=0)$	Units of MeV
E/N	Anomaly ratio in QCD axion models	
$\tilde{C}_{a\gamma\gamma}$	Axion-photon coupling contribution from axion-meson mixing	
$\tan(\beta')$	Ratio of the two Higgs vacuum expectation values	

Table 1. Parameters for the family of axion models. For dimensionful quantities, we also give the units with which they are defined within **GAMBIT**. The first section refers to parameters that are part of the **GeneralALP** model. The second section lists the parameters used in child models of the **GeneralALP**.

<i>Parameter</i>	<i>Value</i>	<i>Comment</i>
Λ_χ	75.5(5) MeV	Ref. [52]
$\tilde{C}_{a\gamma\gamma}$	1.92(4)	Ref. [52]
β	$\hat{\beta} \approx 7.94$	{ Via lattice QCD nuisance likelihood based on ref. [54]
T_χ	$\hat{T}_\chi \approx 147.0 \text{ MeV}$	

Table 2. Nuisance parameters for the QCD axion model. We only quote best-fit estimates of β and T_χ , as the likelihood includes correlations between them (cf. section 3.2).

of the shift symmetry by instanton-like effects is therefore linked to the QCD scale. This connection can be exploited to uniquely determine the parameters $m_{a,0}$, β and T_χ . However, there are uncertainties from theory, experiment and lattice QCD simulations that should be taken into account. We treat these as nuisance parameters; table 2 provides an overview.

Exploiting the connection to QCD, we can replace the parameter $m_{a,0}$ of the **GeneralALP** model with the energy scale Λ_χ , defined via the topological susceptibility at zero temperature, $\Lambda_\chi^4 \equiv \chi(T=0)$, such that

$$m_{a,0} \equiv \frac{\Lambda_\chi^2}{f_a}. \quad (3.4)$$

To determine Λ_χ , we use first principle calculations of the zero-temperature axion mass [52]. We include these results via a Gaussian likelihood

$$\ln(\mathcal{L}) = -\frac{1}{2} \frac{(\Lambda_\chi - \hat{\Lambda}_\chi)^2}{\sigma_{\Lambda_\chi}^2}, \quad (3.5)$$

Figure 2. Nuisance likelihoods for the scale Λ_χ (left, from direct theory calculations [52]) and β and T_χ (right, from lattice QCD [54]). Note that β is correlated with T_χ when the full lattice QCD results are taken into account.

where $\hat{\Lambda}_\chi = 75.5 \text{ MeV}$ and $\sigma_{\Lambda_\chi} = 0.5 \text{ MeV}$ are the most likely values for Λ_χ and its uncertainty from ref. [52].⁷ This is shown in the left panel of figure 2.

The two parameters β and T_χ can then be constrained by using, e.g., the full lattice QCD results from table S.7 in the supplementary material of ref. [54]. We construct a likelihood function by performing a chi-squared fit to the $N = 20$ data points,

$$\ln(\mathcal{L}) = -\frac{1}{2} \sum_{i=1}^N \frac{(X - \hat{X}_i)^2}{\sigma_{X_i}^2}, \quad (3.6)$$

where $X = \log_{10} [\chi(T|\beta, T_\chi)/\chi(T=0)]$ is the logarithm of the normalised topological susceptibility, and $\hat{X}_i$ and $\sigma_{X_i}^2$ are its value and uncertainty for the i th data point.⁸ In figure 2, we show the two-dimensional profile likelihood for β and T_χ . These parameters show a clear correlation. This is expected, as a higher transition temperature implies a steeper slope in the temperature-dependent branch of the axion mass (corresponding to larger β) in order to maintain a good fit to the shape of $m_a(T)$, as determined by lattice QCD. We provide the best-fit values in table 2.

We note that the fit to our functional form for the QCD axion mass (3.3) captures the temperature dependence established by lattice QCD well everywhere except in the region around $T = T_\chi$. This is because (3.3) is not smooth there and the overall fit is poor ($\chi^2 = 55.7$ for 18 d.o.f., which corresponds to a p -value of about 10^{-5}). However, the disagreement stems only from a narrow temperature range, and has no impact on any of our results. Excluding the only data point in that region improves the fit to an acceptable level ($\chi^2 = 21.6$ for 17 degrees of freedom, which corresponds to a p -value of about 0.2).

⁷The authors of ref. [54] find a comparable result of $\Lambda_\chi = 75.6 \pm 1.8(\text{stat}) \pm 0.9(\text{sys}) \text{ MeV}$ using lattice QCD methods.

⁸As the uncertainty quoted in ref. [54] also includes an uncertainty on the value of Λ_χ , we divide the topological susceptibility $\chi = f_a^2 m_a^2$ by their best-fit value for $\chi(T=0)$, and remove this uncertainty by assuming simple error propagation.

We also replace the axion-electron coupling with the model-dependent factor C_{aee}

$$g_{aee} = \frac{m_e}{f_a} C_{aee}, \quad (3.7)$$

and the axion-photon coupling with the model-dependent ratio of the electromagnetic and colour anomalies E/N

$$g_{a\gamma\gamma} = \frac{\alpha_{\text{EM}}}{2\pi f_a} \left(\frac{E}{N} - \tilde{C}_{a\gamma\gamma} \right), \quad (3.8)$$

where $\tilde{C}_{a\gamma\gamma}$ is the model-independent contribution from axion-pion mixing. We use the value obtained in ref. [52] for $\tilde{C}_{a\gamma\gamma}$ and include it as a simple nuisance likelihood,

$$\ln(\mathcal{L}) = -\frac{1}{2} \frac{(X - \hat{X})^2}{\sigma_X^2}, \quad (3.9)$$

with $X = \tilde{C}_{a\gamma\gamma}$ and $\hat{X}$ and σ_X again the most likely value and its uncertainty (see table 2).

Finally, we want to emphasise some subtle considerations about the coupling strengths of the QCD axion model. The possible ratios E/N are rational numbers arising from group theoretical considerations. In this study, we sample over E/N as if it were a continuous parameter, as the possible rational numbers that it can take on are quite densely spaced along the real line, at least over the range of values that we consider. The sometimes so-called “classical axion window” considers a canonical, small and somewhat arbitrarily-defined range of couplings for the prototypical axion models [104, 112–114], arising from only quite a small range in E/N . It has recently been pointed out that the range of choices can, indeed, be extended to include more possibilities [115–118]. To assess the whole range of various axion models, we use the minimum and maximum values for E/N from table IV in ref. [116], so that $E/N \in [-4/3, 524/3]$. These values arise from a systematic study of DFSZ- and KSVZ-type axion models, where the additional heavy quarks in KSVZ-type models have cosmologically safe lifetimes and do not introduce Landau poles below the Planck scale. Furthermore, in DFSZ-type models, the number of Higgs doublets may go up to the maximum of nine.

Note that (3.8) implies the possibility of having $g_{a\gamma\gamma} < 0$ within the valid range for E/N . Note however that all the likelihood functions that we use in this paper only depend on the absolute value of $g_{a\gamma\gamma}$. We therefore plot only $|g_{a\gamma\gamma}|$ in our results, even though we do scan over parameter values that lead to negative couplings in the range $-3.25 \lesssim E/N - \tilde{C}_{a\gamma\gamma} \leq 0$.

3.2.1 KSVZ models

The archetypical axion model is the KSVZ model [41, 42], where the SM is supplemented by one or more electrically neutral, heavy quarks. In our implementation, the KSVZAxion is a child model of the QCDAxion, where the anomaly ratio, E/N , is a free parameter. In these models, axions have no tree-level interactions with fermions. However, there is a loop-induced coupling to electrons due to the axion-photon interaction, which — in the absence of a leading order contribution — must be taken into account [105]:

$$g_{aee} \approx \frac{3\alpha_{\text{EM}}^2}{2\pi} \left[\frac{E}{N} \ln \left(\frac{f_a}{m_e} \right) - \tilde{C}_{a\gamma\gamma} \ln \left(\frac{\tilde{\Lambda}}{m_e} \right) \right] \frac{m_e}{f_a}, \quad (3.10)$$

where $\tilde{\Lambda} \sim 1 \text{ GeV}$ is the QCD confinement scale. Several previous works have used this expression with $\tilde{\Lambda} = 1 \text{ GeV}$ [16, 119], even though this quantity is not uniquely defined. We too assume $\tilde{\Lambda} = 1 \text{ GeV}$, relying on the fact that any deviations enter only logarithmically. Although the anomaly ratio in the original KSVZ model was $E/N = 0$ [41, 42], other assignments are possible. As in ref. [16], we will consider four different KSVZAxion models: $E/N = 0, 2/3, 5/3$, and $8/3$.

3.2.2 DFSZ models

In contrast to the KSVZ model, DFSZ models are obtained by adding an additional Higgs doublet to the SM [43, 44]. This results in direct axion-electron interactions. One can define two manifestations of this model, often called DFSZAxion-I and DFSZAxion-II. The couplings in these two models are given by

$$\begin{aligned} C_{aee} &= \sin^2(\beta')/3, & E/N &= 8/3 & (\text{DFSZAxion-I}) \\ C_{aee} &= [1 - \sin^2(\beta')]/3, & E/N &= 2/3 & (\text{DFSZAxion-II}) \end{aligned} \quad (3.11)$$

where $\tan(\beta')$ is the ratio of the two Higgs vacuum expectation values [105]. Perturbative unitarity requires $0.28 < \tan(\beta') < 140$. It is therefore convenient to replace the parameter C_{aee} in the QCDAxion model by $\tan(\beta')$.

3.3 ALP models with constant mass

As a template model for ALPs that arise as pseudo-Nambu-Goldstone bosons but do not have a temperature-dependent mass, we define the **ConstantMassALP** model. This is mainly for convenience in studies where we want to parametrically explore the coupling space whilst keeping the inverse dependence on f_a , but are not interested in a temperature-dependent mass.

For **ConstantMassALP** models, T_χ is irrelevant because $\beta = 0$, reducing the total number of parameters to five. Similar to QCDAxion models, we replace the ALP mass with a pseudo-Nambu-Goldstone scale Λ and introduce dimensionless coupling constants $C_{a\gamma\gamma}$ and C_{aee} , consistent with the other models:

$$m_{a,0} = \frac{\Lambda^2}{f_a}, \quad g_{a\gamma\gamma} = \frac{\alpha_{\text{EM}}}{2\pi f_a} C_{a\gamma\gamma}, \quad g_{aee} = \frac{m_e}{f_a} C_{aee}. \quad (3.12)$$

4 Observables, experiments and likelihoods

Table 3 gives an overview of the observables and likelihood functions that we use in this paper. In what follows, we give details of the experimental data, computational methods and likelihood implementations that they employ.

4.1 Light-shining-through-wall experiments

Light-shining-through-wall (LSW) experiments shine laser light through a magnetic field onto an opaque material, and attempt to detect it on the other side. A photon may convert into an axion within the magnetic field before the material, pass through it as an axion,

<i>Likelihood/Observable</i>	<i>Comments</i>	<i>References</i>
QCD nuisance parameters	Table 2	[52, 54]; section 3.2
ALPS final limits	vacuum and argon data	[120]; section 4.1
CAST 2007	vacuum data, CCD (2004)	[121]; section 4.2.1 & appendix A
CAST 2017	vacuum data, all detectors	[122]; section 4.2.2 & appendix A
RBF		based on [123, 124]; section 4.3.1
UF		based on [125]; section 4.3.2
ADMX 1998–2009		based on [126–130]; section 4.3.3
ADMX 2018		based on [131]; section 4.3.4
DM relic density	built-in calculator; section 4.4 and appendix B	limits [132]
H.E.S.S.	axion-photon conversion in galactic cluster magnetic fields	based on [133]; section 4.5.1
Supernova limits	axion-photon conversion in magnetic fields of the Milky Way	based on [134]; section 4.5.2
<i>R</i> parameter		[13, 135]; section 4.5.3
<i>Stellar cooling hints (optional likelihood)</i>		
White dwarf cooling hints	only considered in section 5.5	[11, 12, 14, 15]; section 4.5.4

Table 3. Overview of the likelihood functions that we employ in this paper (in the order they are discussed).

and convert back into a photon in the magnetic field on the other side [136–138]. Examples of experiments based on this technique are ALPS [120, 139] and OSQAR [140, 141].

The predicted number of photons on the opposite side of the wall to the laser source is

$$s = \epsilon_{\text{tot}} \mathcal{P}(\gamma \rightarrow a \rightarrow \gamma) \frac{P}{\omega_\gamma} t_{\text{obs}}, \quad (4.1)$$

where ϵ_{tot} is the detector efficiency, P is the laser power, ω_γ the laser energy, and t_{obs} the observation period. $\mathcal{P}(\gamma \rightarrow a \rightarrow \gamma)$ is the probability of a photon converting into an axion and back, in an appropriately aligned magnetic field of length L and strength B . It is given by $\mathcal{P}(\gamma \rightarrow a \rightarrow \gamma) = \mathcal{P}^2(\gamma \rightarrow a)$ with [120]

$$\mathcal{P}(\gamma \rightarrow a) = \mathcal{P}(a \rightarrow \gamma) = \left(\frac{g_{a\gamma\gamma} B L}{2} \right)^2 \text{sinc}^2 \left(\frac{M^2 L}{2\omega_\gamma} \right), \quad (4.2)$$

where $M^2 \equiv m_a^2/2 + \omega_\gamma^2(n-1)$, $\text{sinc}(x) \equiv \sin(x)/x$, m_a is the axion mass, ω_γ is the photon energy, and n is the refractive index of the medium in the experimental setup ($n = 1$ for vacuum).

Our LSW likelihood is based on the final results of the ALPS-I experiment, using data for both evacuated and gas-filled magnets [120]. The ALPS Collaboration took data in “frames” of $t_f = 1$ h each, binning physical pixels of the detector into 3×3 -pixel blocks. ALPs refer to these simply as “pixels” of area $42 \mu\text{m} \times 42 \mu\text{m}$. The collaboration searched their data frames for cosmic rays and signatures of other systematics, over a wide region around the single pixel where most of the laser light would fall in the absence of a wall, referred to as the “signal pixel”.

Figure 3. Validation of our implementation of ALPS result for limits on the axion-photon coupling. We show the 90%, 95% and 99% C.L.s compared to the envelope of their strongest vacuum or gas results. The difference between our likelihood and the limit published by ALPS arises mostly due to the fact that we combine both likelihoods, rather than taking the envelope.

ALPS adjusted the raw ADU values (electron counts) of the signal pixel in an attempt to account for the average background in surrounding pixels. These reduced ADU (**ADUred**) values are obtained by subtracting the average ADU values across all pixels in the region surrounding the signal pixel, from the ADU value of the signal pixel. Doing this for every frame where the laser was on (“signal frames”) and off (“background frames”), ALPS constructed histograms of **ADUred** values for both signal and background. By fitting Gaussian functions to these histograms, they were able to estimate $\hat{b}$ and σ_b , the expected value and standard deviation of **ADUred** for the background, as well as $\hat{o}$ and σ_o , the equivalent quantities for signal frames. An example and more details can be found in figure 2 of ref. [120] (see also ref. [139]).

From (4.1), the expected signal s from photon-axion-photon production per frame of ALPS-I data (with $B = 4.98$ T, $L = 4.2$ m, and $\omega_\gamma = 2.33$ eV) is

$$s \simeq 12.1 \, \epsilon_{\text{tot}} \left(\frac{P}{1096 \text{ W}} \right) \left[\frac{g_{a\gamma\gamma}}{10^{-7} \text{ GeV}^{-1}} \text{sinc} \left(\frac{M^2 L}{2\omega_\gamma} \right) \right]^4. \quad (4.3)$$

Apart from fluctuations in the experimental performance ϵ_{tot} , which amounts to an uncertainty of $\Delta\epsilon_{\text{tot}}/\epsilon_{\text{tot}} \approx 6\%$, the experimental parameters are known to sufficient precision to be fixed to their reference values. We incorporate the uncertainty on ϵ_{tot} into the estimate of the signal prediction, $\sigma_s = s\Delta\epsilon_{\text{tot}}/\epsilon_{\text{tot}}$. As a result of the ADU reduction procedure, the measured signal and background estimators are non-integer numbers. We test the signal-plus-background hypothesis with a Gaussian likelihood:

$$\ln(\mathcal{L}_i) = -\frac{1}{2} \frac{\left[s_i - (\hat{o}_i - \hat{b}_i) \right]^2}{\sigma_{s,i}^2 + \sigma_{\hat{o}-\hat{b},i}^2}. \quad (4.4)$$

We add together the log-likelihoods for data sets $i = 1$ to 3, where two data sets (five and six frames compared to 122 and 47 background frames, respectively) consist of vacuum data and one data set (eight frames compared to 155 background frames) consists of argon gas data.⁹

We show the resulting exclusion limits in figure 3 and compare with the envelope of the strongest vacuum or argon gas limits in ref. [120]. The differences between the published results and our implementation are due to fact that we combine the likelihoods instead of just adopting the more constraining of the two, and also because the authors of ref. [120] used the Feldman-Cousins method [142], assuming Gaussianity and physical signals $s > 0$. Considering the vacuum data alone, the method of ref. [142] gives a slightly stronger limit than our log-likelihood ratio method ($6.5 \times 10^{-8} \text{ GeV}^{-1}$ at 90% confidence limit (C.L.) in the low mass limit, as compared to $6.9 \times 10^{-8} \text{ GeV}^{-1}$ in our implementation). By combining the vacuum and argon likelihoods however, our final limit is somewhat stronger: $g_{a\gamma\gamma} < 5.8 \times 10^{-8} \text{ GeV}^{-1}$ at 90% C.L. at low masses.

4.2 Helioscopes

Axion helioscopes attempt to detect axions produced by interactions in the Sun by observing the solar disc with a “telescope” consisting of a long magnet contained in an opaque casing [143–145]. Solar-produced axions would pass through the casing, convert into photons in the field of the magnet, and be observed in a detector behind the magnet.

Multiple processes in the Sun can produce axions: resonant production in the oscillating electric field of the solar plasma, non-resonant production in solar magnetic fields, and emission from the interaction of electrons with photons, nuclei or one another. The predicted axion-induced photon flux at Earth therefore depends on the assumed solar model, and on the couplings of the axion to both photons and electrons [see e.g. 146].

The dominant process for axion-photon interactions is Primakoff production [147, 148], where photons are resonantly converted into axions in the presence of an atomic nucleus. The rate at which axions of energy E can be produced in a plasma from photons of the same energy is given by [149, 150]

$$\Gamma_{\gamma \rightarrow a} = g_{a\gamma\gamma}^2 \frac{\kappa_s^2 T}{32\pi} \left[\left(1 + \frac{\kappa_s^2}{4E^2} \right) \ln \left(1 + \frac{4E}{\kappa_s^2} \right) - 1 \right]. \quad (4.5)$$

Here, the inverse screening length is given in the Debye-Hückel approximation by

$$\kappa_s^2 = \frac{4\pi\alpha_{\text{EM}}}{T} \left(n_e + \sum_j Z_j^2 n_j \right), \quad (4.6)$$

with n_e denoting the electron number density, and n_j and Z_j representing the number density and charge, respectively, of the j th nucleus. Note that the number densities and temperature, and hence the conversion rate, vary with the radial position r . Using the expression for the axion-photon conversion probability (4.2), the differential photon flux at

⁹Courtesy of Axel Lindner, private communication.

the detector is [121]

$$\frac{d\Phi(E)}{dE} = 2\pi \mathcal{P}(a \rightarrow \gamma) \int_0^{r_s} r \varphi_{\gamma \rightarrow a}(E, r) dr, \quad (4.7)$$

$$\text{where } \varphi_{\gamma \rightarrow a}(E, r) = \frac{R_\odot^3}{2\pi^3 D_\odot^2} \int_r^1 \frac{\rho}{\sqrt{\rho^2 - r^2}} \frac{E \sqrt{E^2 - \omega_{\text{pl}}^2(\rho)}}{e^{E/T(\rho)} - 1} \Gamma_{\gamma \rightarrow a}(E, \rho) d\rho, \quad (4.8)$$

with ρ a dimensionless radial co-ordinate in the Sun, and r a dimensionless radial co-ordinate on the solar disc on the sky. The quantity $D_\odot$ is the (average) Sun-Earth distance, which we take to be one astronomical unit, and $\omega_{\text{pl}}^2(\rho) = 4\pi\alpha_{\text{EM}}n_e(\rho)/m_e$ is the plasma frequency calculated from the electron number density $n_e(\rho)$ and electron mass m_e . The upper limit of outer integral, r_s , controls how much of the inner part of the image of the solar disc on the detector is included in the analysis. This need not always be 1, as the signal-to-noise ratio can be maximised by introducing a cut-off $r_s < 1$ [see e.g. 121].

The contribution to the solar axion flux from axion-electron interactions can be taken into account by including the additional interaction rate $\Gamma_{e \rightarrow a}(E, \rho)$ in (4.8). However, these contributions are not so straightforward to calculate from first principles. This is due to narrow free-bound transition lines [151–154], axion bremsstrahlung [149, 155, 156] and Compton scattering [157–159]. To include these contributions in our signal prediction, we use tabulated data for the axion-electron spectrum provided in ref. [146].

The spectrum in ref. [146] was computed with the 2009 iteration of the AGSS09met solar model [160], which is based on photospheric abundances for non-refractory species and meteoritic abundances for refractory elements [161]. The AGSS09met model is thus the default solar model in GAMBIT, and the one that we use throughout the rest of this paper. We however make use of its latest iteration [162] in preference to the earlier version wherever possible, such as when computing the axion flux from axion-photon interactions. In the limit of small axion mass, the predicted flux from axion-photon interactions deviates by no more than 4.4% between solar models, with the greatest difference occurring between the GS98 [162, 163] and most recent AGSS09met models [162].

Full details of our integration routines for the axion-photon and axion-electron contributions, as well as the options available for the inclusion of solar models for axion physics in GAMBIT, can be found in appendix A.

4.2.1 2007 CAST results

The first of our two CAST likelihoods is based on the CCD results published in 2007 [121] (CCD detector data from the 2004 vacuum run in table 1 of ref. [121]). The other data in ref. [121] are less constraining; including them only improves the upper limit by 1% [121]. We therefore do not provide separate likelihoods for the other runs.

Although the 2007 CAST analysis was based on the solar model of ref. [164], and a follow-up analysis [165] on axion-electron interactions was based on a different model [166], here we use the AGSS09met model of ref. [162]. For both the axion-photon and axion-electron interactions, we integrate the total flux over all 20 energy bins (from 0.8 to 6.8 keV), taking into account the observation time, effective area, and detector efficiency (see ap-

Figure 4. Validation of our implementation of limits arising from the 2007 CAST data [121], including the 2013 re-interpretation in terms of axion-electron couplings [165]. *Left:* exclusion limits for the axion-photon coupling [121]. *Right:* exclusion limits for the product of axion-photon and axion-electron coupling, assuming that axion-electron interactions dominate the axion production inside the Sun. We only make this assumption here to compare our implementation to ref. [165].

pendix A). In this case $r_s \approx 0.23$ for the axion-photon contribution. As we use the interpolated spectrum for the axion-electron contribution (calculated for $r_s = 1$), we can only rescale the resulting flux by an overall factor in order to estimate the flux inside $r_s = 0.23$. We take this number to be 0.826, assuming the findings from the axion-photon interaction also apply in the axion-electron case [121].

Our implementation follows the CAST analysis [121, 165], using a Poisson likelihood in each of the 20 energy bins

$$\ln(\mathcal{L}) = \sum_{i=1}^{20} o_i \ln(s_i + b_i) - \ln(o_i!) - (s_i + b_i), \quad (4.9)$$

where o_i , s_i and b_i are respectively the observed number of photons, the number of expected signal photons, and the expected number of background photons based on observations away from the Sun, in the i th energy bin. In total, 26 photons were observed in the detector during data-taking compared to 30.9 expected background events. The resulting exclusion limits can be found in figure 4.

4.2.2 2017 CAST results

Our implementation of the latest CAST results [122] is analogous to that of the 2007 results, using the signal and expected background counts, exposure and detector efficiencies for the 2017 data.¹⁰

To calculate the signal predictions, we integrate the axion-photon and axion-electron fluxes over each of the 10 energy bins from 2 to 7 keV, and then scale the predictions by the effective exposure for each of the ten datasets in ref. [122].

¹⁰I. Irastorza, private communication.

Figure 5. Validation of our implementation of the 2017 CAST limits [122]. Differences are mainly due to our simplified implementation of the likelihood function, which does not employ event-level information.

The exclusion limits presented in ref. [122] are based on an unbinned likelihood. Here, we treat each energy bin in each of the ten datasets as a separate counting experiment, combining them into a binned Poisson likelihood

$$\ln(\mathcal{L}) = \sum_{j=1}^{12} \sum_{i=1}^{10} o_i^{(j)} \ln(s_i^{(j)} + b_i^{(j)}) + \ln(o_i^{(j)}!) - (s_i^{(j)} + b_i^{(j)}). \quad (4.10)$$

Here $o_i^{(j)}$, $s_i^{(j)}$ and $b_i^{(j)}$ are respectively the observed number of photons, the expected number of signal events and the expected number of background events in the i th energy bin of the j th experiment. In total, 226 photons were observed in the detector during data taking compared to 246.6 expected background events. Our slightly different choice of likelihood function to the original CAST analysis is significantly simpler, because it does not require event-level information — but it still reproduces the published exclusion limits rather well (see figure 5).

4.3 Haloscopes (cavity experiments)

Axion haloscopes are designed to detect DM axions by resonant axion-photon conversion inside a tunable cavity [143, 144]. Microwave cavities are the most sensitive axion experiments in existence, but only cover a small mass range compared to other techniques. The ability of haloscopes to detect axions therefore directly depends on their cosmological abundance, and to a lesser extent, their velocities in the Galactic halo [167]. Here we consider only the case where axions are fully virialised within the halo.

The power expected to be converted from axions to photons in a cavity is [143, 144, 168]

$$P = g_{a\gamma\gamma}^2 B_0^2 C \frac{\rho_{a, \text{local}} V}{m_{a,0}} Q \min\left(1, \frac{Q_a}{Q}\right), \quad (4.11)$$

Figure 6. *Left:* our implementation of haloscope likelihoods compared to the exclusion limits for the RBF [123, 124], UF [125] and ADMX [126–131, 169, 170] experiments. *Right:* magnified details of the latest ADMX results [131].

where C is the form factor (a dimensionless integral over the $\mathbf{E}$ - and $\mathbf{B}$ -field configuration of the cavity), B_0 is magnetic field strength in the cavity, V is its volume, and Q and Q_a are the quality factors of the cavity and axions, respectively. In this context, Q describes the ratio of stored vs dissipated energy of the cavity, while Q_a is proportional to the axion velocity dispersion (just as Q effectively characterises the bandwidth of the cavity).

The signal prediction also depends on the local DM density in axions, which we obtain by rescaling the local DM density as

$$\rho_{a, \text{local}} = \min\left(\frac{\Omega_a}{\Omega_{\text{DM}}}, 1\right) \rho_{\text{DM, local}}. \quad (4.12)$$

Obtaining exclusion limits from cavity experiments is often quite involved, generally requiring simulation of the selection procedure of the detector [e.g. 127]. Without access to this information, we must approximate the underlying likelihood functions based on the publicly available limits and publications.

In the following, we describe our likelihoods for three different haloscope experiments. An overview of the resulting exclusion limits can be found in figure 6.

4.3.1 RBF results

The Rochester-Brookhaven-Fermi (RBF) collaboration performed a search for axions using several cavities [123, 124]. Table I, eq. (26) and figure 14 in ref. [124] provide useful information for approximating their results.

By (4.11), the axion-induced power in a haloscope is proportional to $\rho_{a, \text{local}} g_{a\gamma\gamma}^2$, which we define as the “reduced signal”

$$s \equiv \rho_{a, \text{local}} g_{a\gamma\gamma}^2. \quad (4.13)$$

The remaining factors from (4.11) are effectively constant across all frequency/mass bins and detectors. The signal is expected to occur in a single frequency bin i , which satisfies

$m_{a,0} \in [\omega_i, \omega_{i+1})$. Using this definition, we adopt an ansatz for the likelihood function inspired by eq. (26) of ref. [124]:

$$\ln(\mathcal{L}) = -\frac{1}{2}\Theta(s - a_i) \frac{(s - a_i)^2}{b_i^2}. \quad (4.14)$$

Here a_i is a threshold parameter, b_i effectively corresponds to an expected standard deviation of the reduced signal, and Θ is the Heaviside step function. The threshold values a_i arise because RBF manually inspected all candidate frequencies in their data over a certain significance level. The two parameters are related as $a_i = Nb_i$, where N is the number of standard deviations required for manual inspection of a candidate signal.

Although table I in ref. [124] would allow us to determine b_i , using the central frequency of the bin as well as the 95% C.L. on the coupling strength, the resulting bins are not quite identical to what is shown in figure 14 of the same paper. We therefore determine a_i and b_i for each frequency bin from the limits in figure 15 of that paper, assuming $N = 4$ in all cases (cf. table I in the same paper). This leads to limits in 14 bins ranging from $m_{a,0} = 4.4$ to $10.1 \mu\text{eV}$ and from $m_{a,0} = 11.2$ to $16.2 \mu\text{eV}$.

4.3.2 UF results

While the results from the University of Florida (UF) collaboration [125] could be implemented in the same way as the RBF experiment, the published data do not allow us to infer the threshold parameter a_i for the one mass bin. However, eq. (6) in ref. [125] quotes the “noise power fluctuation”, which we use as a standard deviation $\sigma_P \approx 2.86 \times 10^{-22} \text{ W}$ for the expected power P . We obtain the expected power for each axion model using the information provided in ref. [125], and check that the quoted limit is comparable to the expected sensitivity (which we obtain by assuming that the observed data is equal to the background expectation). The corresponding likelihood function for the single bin from $m_{a,0} = 5.4$ to $5.9 \mu\text{eV}$ and signal s from (4.13) is given by

$$\ln(\mathcal{L}) = -\frac{1}{2} \left(\frac{P(s)}{\sigma_P} \right)^2. \quad (4.15)$$

4.3.3 ADMX results 1998–2009

The procedure used by the ADMX Collaboration for setting limits in the absence of a detection is highly customised for the experiment [127]. Unfortunately, the necessary information for fully implementing their numerous results [126–130, 169] is not available. Similar to the RBF likelihood, we therefore use the reduced signal (4.13) and the following ansatz for the likelihood:

$$\ln(\mathcal{L}) = -\frac{1}{2}\Theta(s/s_i - a) \frac{(s/s_i - a)^2}{b^2}. \quad (4.16)$$

Here s_i is the known limit in the i th frequency/mass bin, and a and b are free parameters that have to be determined by a fit to published exclusion curves. We do this using the exclusion limits at three confidence levels published in ref. [126], and by assuming that the shape of the likelihood function is representative of the shape in all other ADMX bins over the range from $m_{a,0} = 1.90$ to $3.65 \mu\text{eV}$. Doing so results in $a = 0.0131$ and $b = 0.455$.

Figure 7. *Left:* realignment energy density in QCD axions today as a function of mass and initial misalignment angle (fixing β , T_χ , and Λ_χ to their best-fit values for comparison). To guide the eye, we have included dotted and dashed lines to indicate where axions make up certain fractions of the DM relic density. *Right:* band of $m_{a,0}$ - θ_i combinations (from Diver) to get the correct DM density (including β , T_χ , and Λ_χ as nuisance parameters). We show the results of ref. [54] for comparison and also include the hypothetical case of a “temperature-independent QCD axion” (with Λ_χ as a nuisance parameter, but $\beta = 0$ and T_χ irrelevant) as an example.

4.3.4 ADMX results 2018

In their recent publication, the ADMX collaboration increased the sensitivity of their setup, which is now able to rule out some DFSZ-type models [131] in the range $2.66 \leq m_{a,0} \leq 2.81 \mu\text{eV}$. We approximate the likelihood of this result using the 90% C.L. limits in figure 4 of ref. [131], for the Maxwell-Boltzmann velocity distribution (consistent with the model of the halo velocity distribution that we assume for analysing the results of other searches).

Because the experimental setup changed compared to the previous runs, we do not employ the shape parameters from section 4.3.3 for the 2018 dataset. Instead, as with the UF experiment we assume that ADMX saw no signal events, approximating the 2018 likelihood as

$$\ln(\mathcal{L}) = -\frac{1}{2} \frac{s^2}{\sigma_{\text{eff}}^2(m_{a,0}, s)}. \quad (4.17)$$

Here, the effective standard deviation is given by $\sigma_{\text{eff}}^2 = \sigma_{\text{stat}}^2(m_{a,0}) + \sigma_{\text{sys}}^2(s)$. We infer the statistical contribution σ_{stat} by setting the log-likelihood at the observed values of the limits in figure 4 of ref. [131] to that corresponding to a 90% C.L. for one degree of freedom. We read off σ_{stat} at 194 different masses, and interpolate between them linearly for intermediate mass values (ignoring the narrow region from 2.7302 to 2.7307 μeV where the ADMX limits do not apply, due to radio interference [131]). We add the systematic uncertainty of 13% of the signal prediction (quoted in ref. [131]) in quadrature with the statistical uncertainty.

4.4 Dark matter relic density

The realignment mechanism gives an axion contribution to the observed cold DM relic density. We solve (2.9) numerically as a function of temperature to obtain Ω_a , the fraction

of the critical energy density in axions today; details can be found in appendix B. The resulting energy densities for the QCDAxion are shown in figure 7, similar to the presentation in other works [171, e.g.]. We reiterate that in this paper, we do not consider other contributions to the relic density than vacuum realignment.

The relic density of DM is very well constrained by the most recent *Planck* analysis [132]. We employ a Gaussian likelihood with the central value and standard deviation from [132] ($\Omega_{\text{DM}}h^2 = 0.1188$, $\sigma_{\text{exp}} = 0.0010$), combining the experimental uncertainty in quadrature with a further 5% theory uncertainty,¹¹ $\sigma_{\text{theo}} = 0.05 \Omega_a h^2$, to give

$$\ln(\mathcal{L}) = -\frac{1}{2} \frac{(\Omega_a h^2 - \Omega_{\text{DM}} h^2)^2}{\sigma_{\text{exp}}^2 + \sigma_{\text{theo}}^2}. \quad (4.18)$$

GAMBIT offers two options for this likelihood: a detection or an upper limit. These allow us to demand either that axions are responsible for all DM, or only a fraction. For the upper limit, we simply set $\ln(\mathcal{L}) = 0$ for $\Omega_a < \Omega_{\text{DM}}$ in (4.18). Except where we state otherwise, we show results based on the upper limit option.

4.5 Astrophysical probes

Astrophysical systems can provide significant additional constraints on axions, especially the axion-electron and axion-neutron couplings, which are not well constrained by other probes. Due to their weak interactions with matter, axions can efficiently transport energy across large distances in free space, or through stellar matter, thereby influencing stellar structure and evolution.

Intriguingly, a number of astrophysical systems appear to exhibit an unexplained mechanism of energy transport, which might be due to ALPs: white dwarfs display apparently anomalous cooling rates (section 4.5.4) or deviations in the shape of their luminosity function [181, 182] and highly-energetic gamma rays seem to experience significantly less attenuation through intergalactic space than might be expected [17–22]. Unfortunately, the systematic uncertainties associated with these potential hints of new physics are quite difficult to quantify. Nonetheless, if the observations and associated theoretical uncertainties turn out to be robust, ALPs can indeed explain the observed deviations from expectation.

4.5.1 Distortions of gamma-ray spectra

Axions can be generated in galactic or intergalactic magnetic fields, distorting or dimming the spectra of distant sources [183–186]. Several studies have investigated the effects of ALPs on otherwise featureless spectra [187–192], and used the resulting limits to constrain ALP properties [133, 193–196].

¹¹We adopt this entirely heuristic number from the default value in *DarkBit* [25]. The estimated sub-percent numerical systematic uncertainty of our code is smaller, but the theoretical uncertainty due to the possibility of non-standard cosmologies can be much larger [e.g. 172]. Several authors have considered scenarios designed to avoid overproducing axions, including entropy dilution [8, 173–175], inflationary models [176–179], and hidden magnetic monopoles [180].

The probability of photon-axion conversion in a domain of size ℓ filled with a suitably aligned magnetic field B_{eff} and plasma with electron number density n_e is given by [183, 190]

$$p_1(E) \equiv \mathcal{P}(\gamma \rightarrow a) = \frac{1}{1 + \left(\frac{E_{\text{crit}}}{E}\right)^2} \sin^2 \left(\frac{1}{2} g_{a\gamma\gamma} B_{\text{eff}} \ell \sqrt{1 + \left(\frac{E_{\text{crit}}}{E}\right)^2} \right), \quad (4.19)$$

where ω_{pl} and E_{crit} are the plasma frequency and critical energy, respectively:

$$\omega_{\text{pl}} = \frac{4\pi\alpha_{\text{EM}} n_e}{m_e}, \quad E_{\text{crit}} = \frac{1}{2} \frac{|m_a^2 - \omega_{\text{pl}}^2|}{g_{a\gamma\gamma} B_{\text{eff}}}. \quad (4.20)$$

The quantity E_{crit} describes the energy scale at which photons will efficiently convert into axions in the extragalactic magnetic field domains. In the absence of dust or any other photon absorber, after traversing N such domains of size ℓ , the remaining fraction of photons is [184, 185, 190, 197]

$$\mathcal{P}(\gamma \rightarrow \gamma) = \frac{2}{3} + \frac{1}{3} e^{-\frac{3Np_1}{2}}. \quad (4.21)$$

Equations (4.19) and (4.21) reveal that we do not expect to see any effect due to axions for $E \ll E_{\text{crit}}$, because $p_1 \simeq 0$. For a given photon energy E , this happens for large axion masses m_a and small axion-photon couplings $g_{a\gamma\gamma}$.

For $E \gg E_{\text{crit}}$, on the other hand, conversion is very efficient, but the observed spectrum would simply decrease by a constant factor over the entire energy range. In this case it is also not possible to test the axion hypothesis: the expected spectral normalisation of any source is not well constrained and it therefore has to be a free fitting parameter in the analysis.

It is only possible to constrain models where the critical energy lies within the spectral window of the instrument, such that one end of the spectrum is suppressed, but the other is not [191, 192]. Limits from the distortion of gamma-ray spectra are therefore strongest at axion masses that lead to threshold energies similar to the photon energies observed by the experiment. This explains the characteristic shape of the limits (in particular why this method is not sensitive to axion masses below a certain value; see figure 8).

The H.E.S.S. Collaboration applied this technique to the spectrum of the active galactic nucleus PKS 2155-304, using data obtained with their Cherenkov telescope array [133]. Unfortunately, their signal prediction requires Monte Carlo simulation of magnetic field realisations, which are no longer available.¹²

We therefore approximate their likelihood function for the galactic cluster magnetic field using figures 6 and 7 of ref. [133], based on a scheme that we describe in detail in appendix C. The main idea is to use common interpolation methods inside the published exclusion contours, and to extrapolate to likelihood values outside the known contours using a method that mimics the shape of the known iso-likelihood contours and preserves the mathematical properties of the likelihood. Our approximation procedure is of course

¹²P. Brun, private communication.

Figure 8. Our implementation of the H.E.S.S. exclusion limits from gamma-ray spectral distortions of PKS 2155-304, based on the limits quoted in figures 6 and 7 of ref. [133]. Full details of the likelihood construction can be found in appendix C.

somewhat arbitrary. The advantages of our scheme are that it exactly reproduces the known exclusion curves by construction, and that the general likelihood function is well-behaved. The obvious downsides are that we can neither guarantee that the likelihood in the outermost and innermost regions are completely accurate, nor can we extend it to larger couplings than the values shown in figure 7 of ref. [133].

4.5.2 Supernova 1987A

Supernovae are excellent particle laboratories. Supernova 1987A (SN1987A) provides significant constraints on axion parameters, based on the neutrino burst duration [198–201] and axion-photon interaction in magnetic fields external to the supernova [202, 203].

Our likelihood for SN1987A is based on the results from ref. [134]. The authors of that study derived limits based on the absence of a coincident gamma-ray burst from SN1987A, which should have been observed by instruments on board of the Solar Maximum Mission [204] if axions were produced in the explosion and converted to gamma rays in the Galactic magnetic field.

The gamma-ray spectrum of photons with energy E expected at Earth per unit time from axions produced in SN1987A is

$$\frac{d\Phi_{a \rightarrow \gamma}}{dE} = \frac{1}{4\pi d^2} \frac{d\dot{N}}{dE} \mathcal{P}(a \rightarrow \gamma), \quad (4.22)$$

where d is the distance to the supernova, $\mathcal{P}(a \rightarrow \gamma)$ is the conversion probability (4.19), N is the number of axions created in the supernova and $d\dot{N}/dE$ is the axion spectrum at the source, as predicted from a supernova model. To obtain the measured photon fluence at Earth (gamma rays per unit area during the observation), $\mathcal{F}_{a \rightarrow \gamma}$, equation (4.22) has to

Figure 9. Our implementation of SN1987A limits compared to ref. [134] (dashed line).

be integrated over the energy range and time duration of the observation, while modelling the transport of the axions through the Galactic magnetic field.¹³

From figure 6 in ref. [134], we can see that the photon fluence at Earth for a given axion-photon coupling becomes constant below a certain mass scale m_* and rapidly decreases for bigger masses. This is not surprising, given that most axion experiments lose sensitivity at large masses due to the loss of coherence in axion-photon conversion. We therefore make the following ansatz for the fluence,

$$\mathcal{F}_{a \rightarrow \gamma} = F \left(\frac{g_{a\gamma\gamma}}{5.3 \times 10^{-10} \text{ GeV}^{-1}} \right)^4 \begin{cases} 1 & \text{for } m_a \leq m_* \\ \left(\frac{m_*}{m_a} \right)^b & \text{for } m_a > m_*, \end{cases} \quad (4.23)$$

where $F \approx 0.57 \times 10^{-12} \text{ cm}^{-2}$ is the fluence for small axion masses at the reference value of $g_{a\gamma\gamma} \approx 5.3 \times 10^{-10} \text{ GeV}^{-1}$; we obtain this value by integrating eq. (2.11) in ref. [134].

We determine the best-fit values for m_* and the exponent b from the higher-mass region ($m_{a,0} > 6.0 \times 10^{-10} \text{ eV}$) via a least-squares fit to the fluence contour in figure 6 of ref. [134], giving $\hat{m}_* \approx 5.43 \times 10^{-10} \text{ eV}$ and $\hat{b} \approx 4.02$. The value of $\hat{b}$ that we obtain can be qualitatively understood by examining the axion-photon conversion probability given in (4.19), which can be written as $\mathcal{P}(\gamma \rightarrow a) = \sin^2(c\sqrt{x})/x$, where $c = g_{a\gamma\gamma} B_{\text{eff}} \ell / 2$ does not depend on m_a . For axions masses $m_a \gg \omega_{\text{pl}}$ (with $\omega_{\text{pl}} \sim 4 \times 10^{-12} \text{ eV}$), we have $x \propto m_a^4$ and, since the oscillatory part is washed out by the turbulent magnetic fields, the conversion probability is effectively suppressed by a factor of m_a^4 .

¹³The authors of ref. [134] consider two magnetic field models from refs. [205] and [206]. The latter yields weaker limits and is their reference model as well as the basis for our implementation.

The likelihood for $s = \mathcal{F}_{a \rightarrow \gamma}$ in the absence of a photon signal with background fluctuations $\sigma_{\mathcal{F}}^2$ is then given by

$$\ln(\mathcal{L}) = -\frac{1}{2} \frac{s^2}{\sigma_{\mathcal{F}}^2}. \quad (4.24)$$

Figure 9 shows how the limits obtained from our approximation compare to the original reference.

4.5.3 Horizontal branch stars and R parameter

Weakly-interacting particles can influence stellar evolution by providing an additional energy loss mechanism, cooling stars over the course of their evolution [207–209]. The so-called R parameter, $R = N_{\text{HB}}/N_{\text{RGB}}$, is the ratio between the number of Horizontal Branch (HB) stars, N_{HB} , and upper Red Giant Branch (RGB) stars, N_{RGB} , in Galactic globular clusters. Its value depends on the relative time that stars spend on each branch, which is sensitive to the details of stellar evolution and cooling. Axions are expected to be produced in the cores of both types of stars, but would remove heat more efficiently from the cores of HB stars, reducing the time that they spend on the HB and leading to a reduction in R .

Based on a weighted average of a selection of cluster count observations [210], the observed value is $R_{\text{obs}} = 1.39 \pm 0.03$ [135]. The dependence of the predicted value of the R parameter on the properties of axions can be approximated as [13, 135, 211–213]

$$R_{\text{pred}} \approx 7.33 Y - 0.422 - 0.0949 \left(-4.68 + \sqrt{21.9 + 21.1 x_{a\gamma\gamma}} \right) - 0.00533 x_{aee}^2 - 0.0387 \left(-1.23 - 0.138 x_{aee}^{1.5} + \sqrt{1.51 + x_{aee}^2} \right), \quad (4.25)$$

where Y is the Helium abundance, $x_{a\gamma\gamma} \equiv g_{a\gamma\gamma}/10^{-10} \text{ GeV}^{-1}$ and $x_{aee} \equiv g_{aee}/10^{-13}$. The equation above is valid only if axions are sufficiently light compared to the typical temperatures of the stellar interior, which are $T \sim 10^8 \text{ K} \approx 10 \text{ keV}$ [209], i.e. much higher than the axion masses we consider. Our R parameter likelihood is then simply

$$\ln(\mathcal{L}) = -\frac{1}{2} \frac{(R_{\text{pred}} - R_{\text{obs}})^2}{\sigma_{\text{pred}}^2 + \sigma_{\text{obs}}^2}, \quad (4.26)$$

where σ_{pred} and σ_{obs} are the uncertainties of the predicted and observed values.

The authors of ref. [13] adopted a Helium abundance of $Y = 0.255 \pm 0.002$ [214], leading to a predicted value from standard (axion-free) stellar evolution calculations of $R_{\text{theo}} = 1.45 \pm 0.01$, almost 2σ higher than the observed value. We adopt the updated value for such low-metallicity environments of $Y = 0.2515 \pm 0.0017$ [215], leading to $R_{\text{theo}} = 1.42 \pm 0.01$, entirely consistent with the observed R parameter. The effect of the uncertainty of Y on R_{pred} can be estimated according to (4.25), as $\sigma_{\text{pred}} = 7.33 \times 0.0017 \approx 0.012$. A comparison to the exclusion curves in ref. [13] can be found in figure 10.

4.5.4 White dwarf cooling hints

White dwarfs (WDs) are a particularly interesting environment in which to study axion-electron interactions, due to their electron-degenerate cores [10, 216–218]. Current observations can be interpreted as indicating a need for additional cooling in WDs compared

Figure 10. Comparison of our implemented R parameter likelihood with the 1σ and 2σ contours (dashed and dotted blue lines) of ref. [13]. Note that our adopted likelihood leads to limits only, whereas the results of ref. [13] indicate an almost 2σ preference for signal (hence the presence of 1σ upper *and* lower limits in the results that we plot from ref. [13]). The difference is due to the updated He abundance that we employ here.

<i>Object</i>	<i>Type</i>	<i>Mode</i>		2σ C.L. $g_{aee}/10^{-13}$	<i>References</i>
		k	ℓ		
G117-B15A	DA variable	2	1	[3.4, 6.0]	Figure 5 [11]; [219–222]
R548	DA variable	2	1	[0.30, 6.8]	Figure 1 [12]; [221]
L19-2	DA variable	2	1	< 5.1	Figure 5 [14]
PG 1351+489	DB variable	11	1	< 3.6	Figure 5 [15]

Table 4. Overview of available cooling hints for WD variables of spectral types DA and DB. We list the couplings allowed at 2σ confidence (whether limits or intervals). The numerical values necessary for constructing the corresponding likelihoods were kindly provided by A. H. Córscico and T. Battich. In the case of R548, we have two more data points compared to what is shown in figure 1 of ref. [12].

to standard models. The coupling necessary to explain the cooling with axions has been estimated to be $g_{aee} \sim \mathcal{O}(10^{-13})$ [13]. A more recent analysis also considered these cooling hints in a global fitting framework [16]. Whilst the systematics of such analyses are still a matter of debate, and alternative explanations (whether involving BSM physics or not) are certainly still possible, it is intriguing to investigate the impact on axion global fits of including the WD cooling hints. Due to the speculative nature of these hints, we will do so in a separate (alternative) analysis in section 5.5, presented alongside our main global fits.

WDs typically pulsate, allowing the oscillation of their radii and luminosity to be used to probe their internal structure via astroseismology. The periods, Π , of their pulsations decrease with time, with a rate $X = d\Pi/dt$, which can be related to the energy loss in the system.

Figure 11. Overview of the WD likelihoods available in GAMBIT. We show the separate likelihoods (*left*) as well as the combined result (*right*). To guide the eye, horizontal dashed lines indicate the confidence levels. Note that, taken at face value, the combined constraints have the potential to be significant evidence for an additional cooling channel in WDs.

Refs. [11, 12, 14, 15] simulated the evolution of WDs with and without axions, predicting the period decrease $d\Pi/dt$ in each case. For our predictions of WD cooling rates, we interpolate these results and their stated uncertainties, using natural splines. The specific figures and objects from those papers that our implementation is based on are listed in table 4.¹⁴ Note that the plots in refs. [11, 12, 14, 15] show the quantity $m_{a,0} \cos^2(\beta')$ which, for the DFSZAxiom-I axion model they consider, is proportional to the more fundamental parameters C_{aee} or g_{aee} (cf. section 3.2.2). For values of the axion-electron coupling larger than considered in the simulations (i.e. $g_{aee}/10^{-13} > 5.6$ or 8.4), we assign the likelihood corresponding to the largest simulated coupling. This is a conservative assumption, as the disagreement between prediction and observation will in reality only worsen as the coupling increases further (until the WDs become opaque to axions — but this would occur at couplings well beyond what we consider).

For each WD listed in table 4, we use a simple Gaussian likelihood function for the observed ($X_{\text{obs},i}$) and theoretically-expected ($X_{\text{pred},i}$) period decrease, such that our total WD cooling likelihood is

$$\ln(\mathcal{L}) = -\frac{1}{2} \sum_{i=1}^4 \frac{(X_{\text{pred},i} - X_{\text{obs},i})^2}{\sigma_{\text{pred},i}^2 + \sigma_{\text{obs},i}^2}, \quad (4.27)$$

where, again, the predictions and corresponding uncertainties $\sigma_{\text{pred},i}$ are taken from the respective figures in the references listed in table 4 and interpolated via natural splines. The resulting individual and combined likelihoods can be found in figure 11.

We emphasise that the interpretation of WD cooling is subject to a number of assumptions and caveats. Statistical and systematic uncertainties associated with the inputs

¹⁴For L19-2 there is also a measurement for the $k=2, \ell=2$ mode, which results in a stronger preference for $g_{aee} \neq 0$ [14]. However, we choose the $k=2, \ell=1$ mode, consistent with the other DA variable dwarfs.

and algorithms of stellar models were considered by the authors of refs. [11, 12, 14, 15], but a number of other potential issues remain. These include theoretical modelling of the transition from the main sequence to the WD phase, and the accuracy of the observed period decrease of PG 1351+489. Despite these problems, and in contrast to ref. [16], we include the PG 1351+489 system in our discussion. The authors of ref. [16] exclude this object due to its similarity to R548, and the uncertainties associated with R548 being “more conservative”. However, the different estimated uncertainties in these two systems are in fact due to a real physical effect, namely the difference in the influence of trapped vs non-trapped oscillation modes in the two systems. While the latter might give rise to concerns regarding the understanding of different modes (cf. table 4), we do not conclude that the arguments in favour of excluding PG 1351+489 are strong enough to do so.

5 Results and discussion

In this section, we present the central findings from our global fits of various axion models, identifying the most promising regions in parameter space and comparing the various models. We present frequentist and Bayesian results side-by-side, after discussing our choice of priors for the model parameters. Unless stated otherwise, all C.L.s and C.R.s (Bayesian credible regions) are $1\sigma/2\sigma/3\sigma$ (68.27%/95.45%/99.73%), and all C.L.s are two-sided for two degrees of freedom (d.o.f.). Sections 5.2–5.4 do not include WD cooling hints; these are the subject of a dedicated analysis in section 5.5.

5.1 Sampling algorithms and settings

We use the differential evolution sampler Diver [28] to sample the composite likelihood function and T-Walk [28] to sample the posterior distributions. We employ MultiNest [223–225] primarily to compute Bayesian evidences.

We use the sampler settings established in an earlier study by the GAMBIT Collaboration as starting points [28]. For Diver, we generally use a population size (NP) of 2×10^4 and a tolerance (`convthresh`) of 10^{-4} , and turn off the `lambdajDE` optimisation, preferring to use regular `jDE` for its slightly less aggressive optimisation. In addition to combining samples from various runs, where necessary to resolve fine-tuned regions we increase NP to 3×10^4 or 5×10^4 and/or reduce `convthresh` to 10^{-5} . For T-Walk, we use the default settings for 340 or 544 MPI processes, until reaching a tolerance (`sqrtr - 1`) of 0.01 or 0.005. All initialisation YAML files that we use in this study are available on Zenodo [35]; the exact scans for which we use the different settings can be ascertained by examining the input files. Because we use MultiNest primarily for estimating Bayesian evidences, we set the sampling efficiency (`efr`) to 0.3, as recommend for this task [224], and use 2×10^4 live points (`nlive`) with a tolerance (`tol`) of 10^{-4} .

5.2 General ALP models

Starting at the top of the model hierarchy (see figure 1), we first consider the GeneralALP model. This is a phenomenological model; parameter combinations in this model need not correspond to physical models, as their couplings do not depend on the inverse of f_a . The

<i>Model</i>	<i>Parameter range/value</i>		<i>Scale</i>
GeneralALP	f_a [GeV]	$[10^6, 10^{16}]$	log
	$m_{a,0}$ [eV]	$[10^{-10}, 1]$	log
	$g_{a\gamma\gamma}$ [GeV $^{-1}$]	$[10^{-16}, 10^{-8}]$	log
	g_{aee}	$[10^{-22}, 10^{-10}]$	log
	θ_i	$[-3.141\,59, 3.141\,59]$	flat
	β	$[0, 16]$	flat
	T_χ [MeV]	$[10^{-2}, 10^6]$	log
Local DM density	ρ_0 [GeV/cm 3]	$[0.2, 0.8]$	flat

Table 5. Parameter ranges and scaling for GeneralALP models.

main purpose of the GeneralALP is to provide a straightforward, universal connection to observables and to compare to results in the literature.

Parameter ranges. The parameter ranges and scales that we use for GeneralALP models are given in table 5. Because the axion potential is periodic, all normalised field values are equivalent to a value in the interval $(-\pi, \pi]$. For $g_{a\gamma\gamma}$, g_{aee} , and $m_{a,0}$, there is no obvious range to choose; we adopt parameter ranges encompassing values informed by previous studies and phenomenology. Recall that $g_{a\gamma\gamma}$ could be negative but the likelihood functions in the present work only depend its absolute value. We therefore scan only over positive values of $g_{a\gamma\gamma}$, but label plots with $|g_{a\gamma\gamma}|$, to make it explicit that the results are equally valid for the corresponding negative values. For the local DM density ρ_0 , we adopt the same range as in earlier GAMBIT studies [29–32] (and implement the same likelihood function).

The appropriateness of different ranges on f_a , β , and T_χ depends on the fundamental properties of the symmetries and scales of the underlying ALP model. Although there are theoretical arguments for the existence of ALPs from e.g. string theory [38, 226], they do not provide any quantitative guidance.¹⁵ Apart from the likely case that our calculations become meaningless for $f_a \gtrsim m_{\text{Pl}}$, we can only impose $\beta > 0$, as the axion mass should become smaller as the underlying symmetry is restored at higher temperatures. The ranges in table 5 are therefore an attempt to include a variety of cases around the values known or preferred for the QCD axion.

We do not produce Bayesian results for the GeneralALP model, as no strong physical arguments exist for any particular choice of prior on most of its parameters. The only exception is the initial misalignment angle θ_i , due to the causal structure of the early-Universe cosmology mentioned in section 2.3.1.

Frequentist results. We first scan the GeneralALP model assuming ALPs to be all of DM. The resulting limits on the axion-photon coupling (figure 12) are comparable to summary plots elsewhere in the literature [e.g. 114]. However, we would like to stress that unlike overplotted exclusion limits, the exclusion curve in figure 12 arises from a composite

¹⁵See ref. [227] for a recent study, performing a Bayesian analysis of (string-theory-inspired) axion models, similar to our GeneralALP model with $\beta = 0$ and not including coupling strengths as model parameters.

Figure 12. Profile likelihood for GeneralALP models assuming ALPs to be all of dark matter. The constraints that dominate the exclusion contours are CAST, haloscopes (UF, RBF and ADMX), H.E.S.S., the R parameter, and SN 1987A. We used Diver to sample the profile likelihood (interpolated density plot) and a root-finding algorithm with a local optimisation routine for profiling to determine the 2σ C.L. (dashed line). For comparison, we also show the band of QCD axion models that we consider in this paper (blue shaded region; cf. section 5.3) and the discrete choices for E/N that we use for the DFSZ- and KSVZ-type models (yellow lines; dealt with in detail in sections 5.4 and 5.5.2).

Figure 13. Profile likelihoods (from Diver) for GeneralALP models, assuming they provide all of the dark matter in the Universe. We show limits for the axion-photon and axion-electron couplings (*left*; essentially dominated by the R parameter likelihood, cf. figure 10) and for the absolute value of the initial misalignment angle vs f_a (*right*).

likelihood, and profiling takes into account uncertainties in the local DM density (the only relevant nuisance parameter here).

The left panel of figure 13 shows that the joint constraints on the two coupling parameters are essentially dictated by the constraint on the R parameter (figure 10).

Figure 14. Profile likelihoods (from Diver) for specific GeneralALP models that constitute all of the dark matter in the Universe. We compare the allowed regions between QCD-like ALPs with $\beta = 7.94$ and $T_\chi = 147$ MeV (density plot and black contours) and simpler GeneralALP models with $\beta = 0$ (grey contours). Note that in the figure in the left panel, the region to the right of the grey contours is allowed for this model.

As a consequence, the axion-electron equivalent of figure 12 would show that values of $g_{aee} \gtrsim \text{few} \times 10^{-13}$ are excluded across the entire mass range. The right panel of figure 13 shows the possible combinations of f_a and $|\theta_i|$ that allow the GeneralALP to be all of DM. The extent and shape of this region is mostly due to the limited ranges of $m_{a,0}$, β , and T_χ . The axion potential, and therefore the initial energy density in axions, is proportional to $f_a m_{a,0}$. The observed DM abundance can only be achieved if $f_a m_{a,0}$ is large enough, because $|\theta_i| < \pi$. On the other hand, the axion starts to oscillate when $H \sim m_a$ (section 2.3.1). The associated temperature scale depends on $m_{a,0}$, β , and T_χ and sets the amount by which the axion energy density is red-shifted up to the present day. To obtain the correct abundance in axions today while e.g. going to lower values of f_a , the values of $m_{a,0}$, β , or θ_i must be increased or the value of T_χ decreased.

Within our selected parameter ranges, the profile likelihood does not identify preferred regions for T_χ and β . However, as figure 14 shows, different choices of T_χ and β can lead to different behaviours in the profile likelihoods of other quantities. Due to the temperature dependence of the axion energy density (2.11). For example, in figure 14, we choose a generalised QCDaxion-like model (with $\beta = 7.94$ and $T_\chi = 147$ MeV) and compare it to a ConstantMassALP-like model ($\beta = 0$). The different slope of the exclusion region boundary in the left panel of figure 14 is a consequence of how the energy density scales right after the axion field begins to oscillate. For the allowed range of values for $m_{a,0}$, the corresponding band of possible f_a — given a value for the initial misalignment angle — is also different.

5.3 QCD axions

QCD axions are the most well-studied type of axions to date. Unlike in previous studies, here we take into account the uncertainties due to nuisance parameters (see section 3.2) in every part of the analysis. We also consistently scale the local DM density in axions according to their cosmological abundance. This affects the limits on the axion-photon

<i>Model</i>	<i>Parameter range/value</i>		<i>Prior type</i>
QCDAxion	f_a [GeV]	$[10^6, 10^{16}]$	log
	Λ_χ [MeV]	$[73, 78]$	flat
	$\tilde{C}_{a\gamma\gamma}$	$[1.72, 2.12]$	flat
	E/N	$[-1.333\,33, 174.667]$	flat
	C_{aee}	$[10^{-4}, 10^4]$	log
	θ_i	$[-3.141\,59, 3.141\,59]$	flat
	β	$[7.7, 8.2]$	flat
	T_χ [MeV]	$[143, 151]$	flat
Local DM density	ρ_0 [GeV/cm ³]	$[0.2, 0.8]$	flat

Table 6. Prior choices for QCDAxion models.

interaction from haloscope experiments such as ADMX, as the detector signal in (4.13) is proportional to $\rho_{a,\text{local}} g_{a\gamma\gamma}^2$. The limits on $g_{a\gamma\gamma}$ therefore scale with $1/\sqrt{\rho_{a,\text{local}}}$. We again cap the local axion abundance at 100% of the local DM abundance, and penalise models that predict too much DM via the *Planck* likelihood for $\Omega_{\text{DM}} h^2$ (section 4.4).

Prior choices. The priors that we apply to the QCDAxion parameters can be found in table 6. This model imposes a number of relations between the phenomenological parameters of its parent GeneralALP model, which depend on nuisance parameters, i.e. quantities determined by simulations, theory or experiments, which are only known within an appreciable uncertainty. While we are generally not interested in inference on such parameters, their uncertainties can affect results for the actual parameters of interest. The additional nuisance parameters for QCDAxion models are $\tilde{C}_{a\gamma\gamma}$, Λ_χ , β , and T_χ . For $\tilde{C}_{a\gamma\gamma}$ and Λ_χ , the nuisance likelihood is given by a 1D Gaussian for each parameter, whereas the likelihood for β and T_χ takes into account correlations between the two parameters (section 3.2). We choose flat priors from about -5σ to $+5\sigma$ around the respective central values for all four nuisance parameters.

The range of values that we choose for the anomaly ratio, E/N , is inspired by the selection criteria and range established in phenomenological studies of axion models (cf. section 3.2). While the different preferred models presented in ref. [116] form a discrete set, we assume that there is a continuous band of possible axion models, spanning a range from the lowest ($E/N = -4/3$) to the highest ($E/N = 524/3$) possible value of the anomaly ratio. Given that the number of possibilities grows very quickly if we allow for an arbitrary number of new heavy quarks in KSVZ-type models (where, however, $E/N \leq 170/3$), it is not inconceivable that such a band exists. We treat each value inside the band as equally probable before contact with data, employing a flat prior for E/N .¹⁶ Note that

¹⁶The assignment of weights to the different discrete values or to the different parts of the band is not trivial; it becomes complicated quickly if we consider the general QCDAxion family instead of specific DFSZ-type and KSVZ-type models, because the number of additional components (Higgs doublets or heavy quarks) is not fixed. Although it could be argued that all possible values of E/N are equally likely within each class of model, models with more additional particles might be considered more “contrived”, and hence

this necessarily encompasses negative values for $g_{a\gamma\gamma}$, as discussed in section 3.1; whilst this does not impact our likelihoods (which depend only on $|g_{a\gamma\gamma}|$), it does imply an asymmetric effective prior on the two signs of $g_{a\gamma\gamma}$.

Assigning priors to f_a and C_{aee} is more difficult. For f_a , we choose a range that corresponds to our region of interest in mass: from the largest masses allowed by bounds on hot DM, to highly fine-tuned regions with very small masses and f_a somewhat below the Planck scale. The logarithmic prior reflects our ignorance about the scale of new physics, given that the ability of the original QCD axion to solve the Strong CP problem does not depend on the value of f_a . We choose a generous range for C_{aee} , taking a logarithmic prior around values of order unity, which may be considered the most natural value for C_{aee} . Note that the lower end roughly corresponds to the minimum value that can be constrained by the R parameter for the highest QCD Axion masses we consider. Values any lower will be effectively indistinguishable.

Our choice of priors on ρ_0 and θ_i follow the logic from the previous section on GeneralALP models. Because ρ_0 is rather well constrained by data, the choice of log or flat prior has little impact on the final results. For θ_i , the causal structure of the early-Universe cosmology mentioned in section 2.3.1 means that all values of the initial misalignment angle are equally likely, so a flat prior is most appropriate.

Frequentist results. First, let us consider statistical inference on the axion coupling strengths. There are essentially only upper limits on the axion couplings or the associated model parameters. We begin by focusing on the axion-photon interactions, as determined by the anomaly ratio E/N , shown in the upper row of figure 15. In the left panel we impose the relic density constraint as an upper limit, while in the right panel we demand that axions be all of DM. A notable difference between these two assumptions is that haloscopes (UF, RBF and ADMX) only provide strong limits in the latter case.

If axions are not required to be all of DM, the high-mass (low- f_a) region is excluded by the R parameter and CAST likelihoods (cf. figure 12) except at very low values of E/N . If axions constitute all of the DM in the Universe, these constraints are not relevant because the realignment mechanism cannot produce enough DM when $|\theta_i| \leq 3.14159$ and $m_{a,0} \gtrsim 1$ meV (cf. right panel of figure 7), so the high-mass region is excluded.

We also see slightly lower profile likelihood values for masses below about 0.1 peV. This is due to the role of the axion-electron coupling in the R parameter likelihood: while our updated value for the helium abundance reduces the tension between theory and observations, there is still a slight preference for $g_{aee} \neq 0$. For small masses, however, the maximum allowed value for the axion-electron coupling, $C_{aee} \leq 10^4$, is still not large enough to satisfy this small preference.

less probable *a priori*. This is relevant because it particularly affects the higher values of E/N , which cannot be achieved in the simpler models with only one new quark or two Higgs doublets. Creating a probability density function based on how often the values of E/N occur for all the different cases might hence not reflect the *a priori* probability for each version of the QCD Axion model. There are also significant practical challenges to computing all possible values of E/N when the number of quarks becomes very large, as well as in the most general versions of DFSZ-type models.

Figure 15. Profile likelihoods (from Diver) for QCD Axion models with upper limits (*left*) and matching condition (*right*) for the observed DM relic density. The upper and lower panels show the constraints on the anomaly ratio, E/N , and the absolute value of the initial misalignment angle, $|\theta_i|$, respectively.

In the bottom row of figure 15, we show the allowed values for the magnitude of the initial misalignment angle, with and without the assumption that axions constitute all of DM. Due to the influence of the various nuisance parameters and the relic density likelihood, the allowed region in the right panel is not simply a line, but a band of parameter combinations that reproduce the observed DM density within the allowed uncertainties. This panel also illustrates the well-known result that the initial misalignment angle needs to be fine-tuned, i.e. $|\theta_i| \ll 1$, for QCD axion masses of $m_{a,0} \lesssim 0.1 \mu\text{eV}$.¹⁷ We will investigate this issue in more detail below using a Bayesian analysis.

Bayesian results. Breaking the PQ symmetry before inflation effectively results in a single, homogeneous value for the misalignment angle in the entire observable Universe. This gives a physical motivation for choosing a flat prior on θ_i . Parameter regions in which θ_i must be very small to avoid axion overproduction are hence theoretically less

¹⁷Figure 15 contains combined results from multiple Diver runs, designed to properly sample the most fine-tuned regions of the parameter space at low $|\theta_i|$ and $m_{a,0}$.

Figure 16. Marginalised posteriors (from T-Walk) for QCD Axion models with upper limits (*left*) and matching condition (*right*) for the observed DM relic density. The upper and lower panels show the constraints on the anomaly ratio, E/N , and the absolute value of the initial misalignment angle, $|\theta_i|$, respectively.

appealing. In a Bayesian analysis, we can see and quantify these fine-tuning issues in the (marginalised) posterior distributions, which quantify the degree of belief in certain values of the parameters given data and prior information.

We show marginalised posteriors for the QCD Axion model in figure 16, once again without (*left*) and with (*right*) the requirement that QCD axions are all of DM. As a consequence of fine-tuning in θ_i , the low-mass (high- f_a) region of the parameter space in figure 16 is disfavoured, even when taking the DM relic density as an upper limit only. This is because in the low-mass region, large absolute values of the initial misalignment angles have a small likelihood. An $\mathcal{O}(1)$ value for the magnitude of the initial misalignment angle is *a priori* more probable than finding a value close to zero, due to the flat prior. This conflict leads to fine-tuning becoming increasingly necessary as the axion mass decreases, which is penalised in the Bayesian analysis. Although such parameter combinations might still give valid solutions that evade all constraints, they are not as probable as others.

A similar logic applies to large axion-photon coupling, i.e. large E/N . Due to the fine-tuning in E/N necessary to evade the helioscope and R parameter constraints at large axion

Figure 17. Profile likelihoods (from Diver, *left*) and marginalised posteriors (from T-Walk, *right*) for the mass in QCD Axion models with upper limits (red shading) and matching condition (blue shading) for the DM relic density. The prior-dependence of the marginalised posteriors is investigated in appendix D.

mass (cf. the corresponding profile likelihood in the top left panel of figure 15), the large- $m_{a,0}$ (low- f_a) region in the top left panel of figure 16 is disfavoured in the Bayesian posterior.

If we demand that axions explain all of DM, the consequences are even more dramatic. The most probable axion models are confined to the narrow band in $m_{a,0}$, visible in the upper right panel of figure 16. This mass range presents a feasible target for haloscope searches, and ADMX in particular is already beginning to cut into these models from the left (low-mass end). This also explains why the band of $m_{a,0}$ - θ_i values in the bottom right panel of figure 16 is not continuous, but disrupted around two points. These correspond to the ADMX and RBF/UF haloscope searches, respectively. While the RBF and UF haloscopes cannot reach as far down into the coupling space as ADMX, they do still constrain a significant fraction of the coupling range.

The fact that the Bayesian analysis singles out a well-defined range for the QCD axion mass becomes even more apparent in figure 17, where we compare one-dimensional profile likelihoods and marginalised posteriors for the axion mass. The frequentist approach does not yield a clear preference for any mass range, but the posterior distributions are strongly peaked around $m_{a,0} \sim 100 \mu\text{eV}$. Clearly, such a result is not completely prior-independent, and we discuss the impact of adopting different priors in appendix D. Nevertheless, it is appealing that a Bayesian analysis can identify a preferred region of $m_{a,0}$ which, intriguingly, falls into the range that can be covered by experimental searches. Indeed, the impact of ADMX and other haloscopes already manifests itself as dips in the right panel of figure 17.

The marginalised posterior in figure 17 allows us to infer a preferred QCD Axion mass range. When demanding that axions explain all of DM, we find that the 95% equal-tailed credible interval for the axion mass is $0.12 \mu\text{eV} \leq m_{a,0} \leq 0.15 \text{ meV}$; allowing them to constitute a fraction of DM, this becomes $0.48 \mu\text{eV} \leq m_{a,0} \leq 3.8 \text{ meV}$. These numbers have minimal dependence on the adopted prior for E/N , but a stronger dependence on the choice of priors for C_{aee} and f_a (appendix D).

Figure 18. Profile likelihoods (from Diver, *left*) and marginalised posteriors (from T-Walk, *right*) for $\Omega_a h^2$ in QCD Axion models with upper limits (red shading) and matching condition (blue shading) for the DM relic density.

We also note that if the PQ symmetry is broken after inflation, the preferred axion mass range will generally shift to larger values due to averaging of the energy density and inclusion of topological defects (cf. section 2.3). The lower bounds on $m_{a,0}$ that we quote can therefore be viewed as robust against changes of assumptions about inflation.

Finally, we also show the one-dimensional profile likelihoods and marginalised posteriors for the QCD Axion relic density in figure 18. Demanding that axions be all of the DM effectively dominates the outcome of this analysis. Using the DM relic density as an upper limit causes the profile likelihood to essentially follow the relic density likelihood function (left panel). In a Bayesian analysis, however, we immediately see that QCD Axions are not expected to generally provide all of the DM in the Universe, given our definition of the parameter space and priors. Imposing the DM relic density as an upper limit, the median axion relic density is 6.5×10^{-3} , or about 5% of the observed DM abundance. The 95% credibility equal-tailed preferred range is $6.8 \times 10^{-6} \leq \Omega_a h^2 \leq 0.10$, which corresponds to between about 0.006% and 90% of the cosmological density of DM. This demonstrates that in the pre-inflationary PQ symmetry-breaking scenario, although QCD Axions can provide a sizeable contribution to the DM density of the Universe, they probably do not contribute all of DM. Again, we stress that these statements are sensitive to the adopted prior on C_{aee} and f_a (appendix D).

5.4 DFSZ- and KSVZ-type models

The DFSZ-type (DFSZ Axion-I, DFSZ Axion-II) and KSVZ-type (KSVZ Axion) models differ from their parent model, the QCD Axion, in that they specify the axion-photon and axion-electron coupling strengths, or at least limit them to a well-defined range for a given axion mass. They are but a few of the many possible phenomenologically-inspired models, but they serve as interesting archetypes of their respective subclasses to compare with more general QCD Axion models.

<i>Model</i>	<i>Parameter range/value</i>		<i>Prior type</i>	<i>Comments</i>
	f_a [GeV]	$[10^6, 10^{16}]$	log	Applies to all
	Λ_χ [MeV]	[73, 78]	flat	Applies to all
	$\tilde{C}_{a\gamma\gamma}$	[1.72, 2.12]	flat	Applies to all
	θ_i	$[-3.141\,59, 3.141\,59]$	flat	Applies to all
	β	[7.7, 8.2]	flat	Applies to all
	T_χ [MeV]	[143, 151]	flat	Applies to all
DFSZAxion-I	E/N	8/3	delta	
	$\tan(\beta')$	[0.28, 140.0]	flat	
DFSZAxion-II	E/N	2/3	delta	
	$\tan(\beta')$	[0.28, 140.0]	log	
KSVZAxion	E/N	0, 2/3, 5/3, 8/3	delta	Various discrete choices
Local DM density	ρ_0 [GeV/cm ³]	[0.2, 0.8]	flat	

Table 7. Prior choices for DFSZAxion-I, DFSZAxion-II and KSVZAxion models. Note that the priors listed in the first section of the table apply to all three models.

Prior choices. Our prior choices for the DFSZ and KSVZ model can be found in table 7. For most of them, the rationale is the same as for QCDAxions presented in section 5.3. The only differences are in the parameters related to couplings. We fix E/N to some typical values considered previously in the literature [e.g. 16]. The range that we choose for $\tan(\beta')$ in DFSZ-type models reflects the values allowed by perturbativity bounds [228]. Our choice of a log prior for $\tan(\beta')$ reflects the assumption that each possible Higgs vacuum expectation value is equally likely; indeed, any sensible prior choice for this parameter should reflect the fact that the two Higgs doublets may be interchanged, and the prior should be invariant under inversion of the ratio of vacuum expectation values.

Frequentist results. The DFSZ- and KSVZ-type models are essentially restrictions of the allowed QCD axion couplings. We therefore only expect to see qualitative differences in the results from the different models where the DFSZ and KSVZ interaction strengths cannot be tuned sufficiently to evade constraints from haloscopes (if axions make up all of DM) and the R parameter.

Figure 19 shows the profile likelihood constraints on various axion models, imposing the DM density as an upper limit. We can see that the upper limit on the axion mass depends on the value of E/N in a given model (cf. figure 12), giving KSVZAxion models with $E/N = 5/3$ the largest allowed parameter space out of all the models compared here. The different values for E/N are also the reason for the different positions of the peaks in $m_{a,0}$; the slight preference for non-zero couplings in the R parameter likelihood requires slightly different axion masses for different E/N .

Demanding that axions explain all of DM, figure 19 would change slightly. All models with $m_{a,0} \gtrsim 0.1$ meV would be ruled out (not being able to provide all of the DM through the realignment mechanism), and ADMX would partially constrain all models except those with $E/N = 5/3$ (cf. figure 12).

The relation between the DFSZ- and KSVZ-type models and their parent QCDAxion

Figure 19. Profile likelihood (from Diver) for the axion mass in the different DFSZ- and KSVZ-type models, imposing the DM relic density as an upper limit.

Figure 20. Profile likelihood (from Diver) for KSVZAxion (*left*), DFSZAxion-I (*right*; blue contours), and DFSZAxion-II models (*right*; coloured region and black contours), compared to the profile likelihood for the QCDAxion model (grey lines). All results use the observed relic density of DM as an upper limit on the relic axion abundance.

model determines their allowed axion-electron couplings. Figure 20 shows how the allowed parameter space in the $m_{a,0}$ - C_{aee} parameter plane of the QCDAxion model is constrained further by imposing additional relations between the different model parameters in the KSVZ- and DFSZ-type models.

This is most striking in the case of the KSVZ-type models, for which C_{aee} is only induced at the loop level and depends directly on $m_{a,0}$ (3.10). Note that the ordering of the KSVZAxion regions is also non-monotonic in E/N due to the difference term in (3.10). The finite sizes of the allowed parameter regions are simply a result of the nuisance parameters included in the relation between C_{aee} and $m_{a,0}$. For DFSZ-type models, C_{aee} depends on the additional parameter $\tan(\beta')$, which makes it possible to accommodate a wide range of axion-electron couplings. However, the parameter space is also more constrained in this case, as very large values of C_{aee} cannot be realised given other constraints on $\tan(\beta')$.

Also note that, due to the different coupling structure in DFSZAxion-I and DFSZAxion-II models (3.11), the same range for $\tan(\beta')$ translates into a lower minimal value of C_{aee} in DFSZAxion-II models than in DFSZAxion-I models. The resulting possible range for C_{aee} in DFSZAxion-II models (and KSVZAxion models with $E/N = 5/3$) also extends to slightly lower values than the prior box that we chose for QCDAxions.

Bayesian results. We use the nested sampling package MultiNest to estimate the Bayesian evidences $\mathcal{Z}(\mathcal{M})$ for each model $\mathcal{M}$. From these we construct the Bayes factor [229–231]

$$\mathcal{B} \equiv \frac{\mathcal{Z}(\mathcal{M}_1)}{\mathcal{Z}(\mathcal{M}_2)} \equiv \frac{\int \mathcal{L}(\text{data} | \boldsymbol{\theta}_1) \mathcal{P}_1(\boldsymbol{\theta}_1) d\boldsymbol{\theta}_1}{\int \mathcal{L}(\text{data} | \boldsymbol{\theta}_2) \mathcal{P}_2(\boldsymbol{\theta}_2) d\boldsymbol{\theta}_2}, \quad (5.1)$$

where $\mathcal{M}_1$ and $\mathcal{M}_2$ are the two models under investigation, $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$ are their parameters, $\mathcal{P}_1$ and $\mathcal{P}_2$ are their priors and $\mathcal{L}$ is the likelihood. The Bayes factor is connected to the odds, i.e. the ratio of posterior probabilities, of the models being correct:

$$\frac{\mathcal{P}(\mathcal{M}_1 | \text{data})}{\mathcal{P}(\mathcal{M}_2 | \text{data})} = \mathcal{B} \frac{\mathcal{P}(\mathcal{M}_1)}{\mathcal{P}(\mathcal{M}_2)}. \quad (5.2)$$

In this paper, we assign equal prior probabilities to both models being correct, i.e. choose $\mathcal{P}(\mathcal{M}_1)/\mathcal{P}(\mathcal{M}_2) = 1$, so the Bayes factor is the same as the posterior odds ratio.

Using MultiNest’s nested sampling (as opposed to importance nested sampling) estimates for evidences, we calculated the odds in favour of KSVZ- and DFSZ-type axion models over the QCDAxion model. In terms of the commonly used scale for Bayes factors [229, 230], we find that there is generally no noticeable evidence for or against any of these models, which would require an odds ratio of more than 3:1 (or less than 1:3).

Imposing the relic DM density as an upper limit, the odds in favour of any KSVZ- or DFSZ-type axion models, compared to QCDAxions, are 2:1. If we demand that axions constitute all of DM, the odds reduce to 1:1.

The outcome of the model comparison is not surprising, as we have not included any positive evidence for axions at this stage. We will discuss in the following section how these conclusions change when including WD cooling hints.

5.5 Cooling hints

Observables related to stellar cooling offer a unique opportunity to constrain the axion parameter space. If future observations confirm the need for additional cooling channels to explain the observed decreases in WD pulsation periods, we may be able to use WDs to measure the axion mass and coupling strengths. In this section we add the likelihoods related to the WD cooling hints to our analysis, emphasising once again the caveats and difficulties associated with assigning uncertainties to the model predictions (cf. section 4.5.4). Here our prior choices for each model are the same as in the preceding sections. A detailed numerical comparison of our results to previous works [13, 16] is not meaningful due to differences in the choice of WD likelihood function, but the findings are qualitatively similar.

Figure 21. Profile likelihoods (from Diver) for QCD Axion models with upper limits (*left*) and matching condition (*right*) for the observed DM relic density and including cooling hints. The upper and lower panels show the constraints on the anomaly ratio, E/N , and for the absolute value of the initial misalignment angle, $|\theta_i|$, respectively.

5.5.1 QCD axions

Previous studies have mostly considered the phenomenological couplings $g_{a\gamma\gamma}$ and g_{aee} or specific QCD axion models with fixed E/N . Here, we investigate which parts of the broader QCD Axion parameter space can explain the cooling hints.

Frequentist results. Figure 21 is the cooling-hint equivalent of figure 15, summarising the allowed anomaly ratio and magnitude of the initial misalignment angle. The only notable difference is at $m_{a,0} \lesssim 0.1 \mu\text{eV}$ ($f_a \gtrsim 3 \times 10^{13} \text{ GeV}$), where none of the possible values for C_{aee} under consideration is large enough to fully account for the anomalous cooling. Recall that $g_{aee} \propto C_{aee} m_{a,0}$ (3.7) and that the cooling hints point towards a relatively narrow range of couplings g_{aee} (figure 11). The overall effect of the cooling hints is therefore to disfavour lower masses. Had we chosen the range of possible values for C_{aee} to be smaller, these constraints would extend to even larger values of $m_{a,0}$ (and vice versa if we had permitted even larger values of C_{aee}).

The right panels of figure 21 show that QCDAxion models can satisfy the cooling hints *and* be all of the DM in the interval $0.1 \lesssim m_{a,0}/\mu\text{eV} \lesssim 300$. The lower bound on this mass region depends on the largest allowed value for C_{aee} .

It is interesting to consider how good the fit of the QCDAxion model is in an absolute sense. Most constraints are easily satisfied by the best-fit point, such that the corresponding partial likelihoods give p -values of order one, which we will not discuss further.¹⁸ One exception is the fit to the temperature dependence of the QCD axion mass, which gives a p -value of order 10^{-5} (see section 3.2). Ignoring this likelihood (and the two model parameters constrained by it) we are left with 7 model d.o.f. and 48 data d.o.f. when including the WD cooling hints; without the cooling hints, the data d.o.f. is 43. The corresponding p -value is 0.30 with cooling hints included, and 0.60 without. The decrease in p -value when including the WD cooling hints results from the slight discrepancies between the cooling hints themselves (cf. figure 11) and their slight tension with the R parameter likelihood.

Bayesian results. Selected results from the Bayesian analysis of QCD axions in combination with cooling hints can be found in figure 22. Compared to the Bayesian results without cooling hints in figure 16, we can see that the preferred mass regions in the $m_{a,0}$ - θ_i plane get narrowed down slightly when we impose the DM relic density constraints as an upper limit (bottom left plot). However, for the anomaly ratio, E/N , this is not the case (top left plot). Generally speaking, these results identify the most credible regions for a compromise between QCD axions fitting the cooling hints and “naturally” not overproducing DM (which prefers masses of $\mathcal{O}(1 \text{ to } 100 \mu\text{eV})$, cf. figure 16). Despite the slight differences, which might also depend on the adopted priors, the overall results with and without cooling hints are remarkably similar. This is mainly due to the influence of the R parameter likelihood included in both cases and its slight preference for non-zero couplings.

The influence of the cooling hints is illustrated further in figure 23, which shows the regions of highest posterior probability in the $m_{a,0}$ - C_{aee} parameter plane with and without the inclusion of cooling hints. Since the cooling hints strongly require $g_{aee} \sim 3 \times 10^{-13}$ (cf. right panel of figure 11), we find the highest posterior probabilities along a line of constant $C_{aee} m_{a,0}$. The chosen range of C_{aee} then implies that the cooling hints can only be explained for $m_{a,0} \gtrsim 0.3 \mu\text{eV}$. For values of $C_{aee} \sim 1$, the cooling hints would point towards meV-scale axions, which is incompatible with the requirement $\Omega_a \sim \Omega_{\text{DM}}$. This regime is therefore disfavoured in the right panel of figure 23. Not including the cooling hints essentially only results in an upper limit in the most credible mass regions close to where the line of constant g_{aee} was (grey contours in figure 23), due to the R parameter likelihood.

As mentioned before, QCD axions can account for both the cooling hints and all of the DM in the Universe (cf. figures 21 and 22). However, because the posterior probability in figure 22 is normalised, one cannot infer from this plot if these solutions occur naturally or if considerable fine-tuning is required. Figure 24 gives an idea of the “naturalness” of QCDAxion DM by showing the marginalised posterior as a function of $m_{a,0}$ and $\Omega_a h^2$ (with and without the WD cooling hints). In the colour density plots of both panels

¹⁸The exact numerical values depend on how many data and model d.o.f. one takes into account, which is often ambiguous.

Figure 22. Marginalised posteriors (from T-Walk) for QCDAxion models with upper limits (*left*) and matching condition (*right*) for the observed DM relic density and including cooling hints. The upper and lower panels show the constraints on the anomaly ratio, E/N , and for the absolute value of the initial misalignment angle, $|\theta_i|$, respectively.

Figure 23. Marginalised posteriors (from T-Walk) for QCDAxion models with upper limits (*left*) and matching condition (*right*) for the observed DM relic density. We show the constraints on the axion-electron coupling, C_{aee} , without (grey lines) and with (black lines and coloured regions) the inclusion of cooling hints.

Figure 24. Marginalised posteriors (from T-Walk) for QCD Axion models with upper limits (density plots and black contour lines) and matching condition (grey contour lines) for the observed DM relic density. We show the constraints on the energy density in axions today, $\Omega_a h^2$, without (*left*) and with (*right*) the inclusion of cooling hints.

in figure 24, we can see that the scan finds credible parts of the parameter space where axions account for a sizeable fraction of the DM while being consistent with all experiments and observations. The differences between including and not including the cooling hints regarding the preferred regions of $\Omega_a h^2$ are rather small, consistent with the other plots.

Similar to the discussion at the end of section 5.3, we can infer the most credible regions for the relic abundance of axions as well as for $m_{a,0}$. The preferred axion mass is very similar with or without the inclusion of the cooling hints. Imposing the DM relic density as an upper limit, we find $0.70 \mu\text{eV} \leq m_{a,0} \leq 2.8 \text{ meV}$ at 95% credibility (equal-tailed interval); demanding that axions be all of the DM, this becomes $0.41 \mu\text{eV} \leq m_{a,0} \leq 0.14 \text{ meV}$. Including the cooling hints slightly modifies the preferred range for $\Omega_a h^2$. At 95% credibility (equal tails), $1.0 \times 10^{-5} \leq \Omega_a h^2 \leq 0.10$, corresponding to 0.009–90% of DM. The median value is $\Omega_a h^2 = 9.3 \times 10^{-3}$, or about 8% of the observed DM.

Finally, let us return to the $m_{a,0}$ – $g_{a\gamma\gamma}$ parameter plane, as discussed in the GeneralALP model without cooling hints (see figure 12). In figure 25, we contrast the naïve bounds on the parameter space (from the phenomenological constraints on GeneralALP models and the maximum possible value of E/N) with the regions preferred by a Bayesian analysis. These regions are not only determined by the constraints from data (satisfying the cooling hints in both panels and matching the DM density in the right panel), but also by the fine-tuning in some parts of the parameter space. Fine tuning is necessary for avoiding overproduction of DM at small $m_{a,0}$, and for achieving low values of $g_{a\gamma\gamma}$ through cancellations between E/N and $\tilde{C}_{a\gamma\gamma}$ at large $m_{a,0}$ (cf. 3.8). For our adopted priors, the most credible parameter regions correspond to a few orders of magnitude around $m_{a,0} \sim 10 \mu\text{eV}$ and $g_{a\gamma\gamma} \sim 10^{-12} \text{ GeV}^{-1}$. In appendix D we discuss how choosing different priors may affect these conclusions.

Figure 25. Marginalised posteriors (from T-Walk) for QCDAxion models with upper limits for the observed DM relic density (*left*) and matching condition (*right*). We show the constraints on the absolute value of the axion-photon coupling, $|g_{a\gamma\gamma}|$, together with models *outside* of the band of axion models (red line and shading) and the (frequentist) 2σ C.L. (dashed lines) for comparison. The prior-dependence of these results is investigated in appendix D.

Figure 26. Profile likelihoods (from Diver) for DFSZAxion-I (*left*) and DFSZAxion-II (*right*) models, compared to the profile likelihood for the QCDAxion model (grey lines), taking the observed DM abundance as an upper limit on the number of axions. Contours show the 1σ and 2σ confidence regions.

5.5.2 DFSZ- and KSVZ-type axions

DFSZ-type models have intrinsically larger coupling to electrons than KSVZ-type models, which only obtain their interactions with electrons at loop level. DFSZ models are therefore the natural choice to account for the potential WD cooling anomalies [16], by way of an axion-electron coupling of $g_{aee} \sim 3 \times 10^{-13}$.

Frequentist results. Figure 26 shows the profile likelihood for the DFSZAxion-I and DFSZAxion-II models, compared to the band that maximises the profile likelihood for the QCDAxion model. Clearly both DFSZ-type models can accommodate the cooling hints with large axion-electron couplings. We already saw in figure 20 that even without the

<i>Model</i>	<i>E/N</i>	$\hat{m}_a$ meV	$\hat{C}_{a\gamma\gamma}$	$\hat{g}_{a\gamma\gamma}$ $10^{-12} \text{ GeV}^{-1}$	$\hat{C}_{aee}$ 10^{-3}	$\hat{g}_{aee}$ 10^{-15}	$-2\Delta \ln(\hat{\mathcal{L}})$
DFSZAxion-I	8/3	8.77	0.746	1.33	333	262	0.71
DFSZAxion-II	2/3	9.19	1.25	2.35	309	254	1.29
KSVZAxion	0	15.4	1.92	6.04	0.118	0.163	9.58
	2/3	24.2	1.25	6.18	0.0272	0.0589	9.58
	5/3	24.2	1.25	6.18	0.0272	0.0589	9.58
	8/3	40.8	0.747	6.21	0.451	1.65	9.55

Table 8. Best-fit values for DFSZAxion-I, DFSZAxion-II, and KSVZAxion models when imposing an upper limit on the DM relic density, $\Omega_a \lesssim \Omega_{\text{DM}}$. In the final column, we compare the likelihood of the respective best-fit points to QCDAxion models. Note that the QCDAxion model has two more degrees of freedom than the KSVZ models, and one more than the DFSZ-type models.

cooling hint likelihood, the KSVZAxion fails to achieve large C_{aee} values; with cooling hints included, the highest likelihood regions for KSVZAxion models are therefore essentially the same as in figure 20.

Table 8 gives the best-fit values for the six classic axion models that we consider in this section, under the requirement that they do not exceed the observed abundance of DM. We do not report best-fit values for $|\theta_1|$, as even with the maximum value included in our scans, axions only account for a few percent of the observed DM abundance. For each model, we calculate $\Delta \ln(\hat{\mathcal{L}})$, the logarithm of the ratio of the best-fit likelihood relative to the QCDAxion model. As anticipated, DFSZ-type models perform much better than KSVZAxions (at the expense of having one additional degree of freedom). This is because DFSZ-type models can easily reach the required axion-electron coupling to fit the cooling hints with masses of order $m_{a,0} \sim 10 \text{ meV}$ (as noted previously [16]). For KSVZ-type axions, the masses required to naturally fit the cooling hints are about three to four orders of magnitude larger (see e.g. 3.10), and the associated axion-photon coupling is therefore in conflict with the R parameter likelihood (as well as hot DM bounds, which are not included). Nevertheless, even for KSVZAxion models there is still a slight preference compared to having no axion at all, which corresponds to $-2\Delta \ln(\hat{\mathcal{L}}) = 10.54$.

For DFSZ-type models, the DFSZAxion-I model gives a better fit than the DFSZAxion-II model. This is due to the influence of the R parameter likelihood, which combines with the cooling hints to force the axion-photon coupling to $g_{a\gamma\gamma} \lesssim 2 \times 10^{-11} \text{ GeV}^{-1}$ (at 95% C.L.). The best-fit point is therefore a balance between reaching high enough g_{aee} and minimising $g_{a\gamma\gamma}$. The maximum value of C_{aee} for DFSZAxion-I models is about a factor of 1.08 larger than for DFSZAxion-II models, due to perturbativity constraints on $\tan(\beta')$, whereas the axion-photon couplings are about a factor of 0.6 lower, yielding a better fit to the R parameter likelihood at any given mass.

Next, let us briefly consider the case where we demand that the classic axion models provide all of the DM. This results in much poorer maximum likelihood values of around $-2\Delta \ln(\hat{\mathcal{L}}) \approx 10.5$ compared to QCDAxion models for *all* KSVZ- and DFSZ-type models that we consider in this paper. This is not surprising because, unlike some other QCDAxion

Figure 27. Marginalised posteriors (from T-Walk) for DFSZAxiom-I (*left*) and DFSZAxiom-II (*right*) models, compared to the marginalised posterior for the QCDAXiom model (grey lines), taking the observed DM relic density as an upper limit on the axion relic density and including cooling hints. We show only the 68.27% and 95.45% C.R.s for all models.

models, they cannot account for both the cooling anomalies and DM.¹⁹ The maximum-likelihood regions are also highly degenerate in this case, as none of the fits is actually “good” (with a maximum likelihood comparable to the case without any axion).

Finally, thanks to the model hierarchy shown in figure 1, we can perform nested hypothesis tests to determine whether or not the more constrained, specific DFSZAxiom and KSVZAxiom models are disfavoured compared to the broader class of QCDAXiom models. Without the cooling hints, QCDAXiom, DFSZAxiom and KSVZAxiom models cannot be discriminated. This situation changes if the WD cooling hints are taken into account. Imposing the measured DM density as an upper limit only, KSVZAxiom models (null hypothesis) can be rejected with respect to QCDAXiom models (alternative hypothesis) with a p -value of 8.3×10^{-3} . DFSZAxiom models as a null hypothesis, on the other hand, cannot be rejected ($p \geq 0.26$). If we instead demand that axions be all of DM, both DFSZAxiom and KSVZAxiom models can be rejected with respect to the QCDAXiom model with p -values of 1.2×10^{-3} and 5.3×10^{-3} , respectively.

Bayesian results. Figure 27 shows posteriors for the DFSZ-type models compared to the QCDAXiom, in the $m_{a,0}$ - C_{aee} parameter plane (see also figure 23). These regions resemble the ones found in figure 26. The log prior on $\tan \beta'$ enables both models to achieve sufficiently large values of C_{aee} quite naturally, despite their structural differences. At first sight, DFSZAxiom-II models appear to be able to occupy a larger region of parameter space in the 68.27% C.R. However, this is mainly a reflection of the effective prior on C_{aee} , which for the DFSZAxiom-I model is strongly peaked towards the upper boundary of the accessible parameter space in figure 27. In contrast, the effective prior in C_{aee} is almost flat in the DFSZAxiom-II model. In terms of the fundamental parameters (i.e. $\tan \beta'$), the credible

¹⁹Note that $\theta_i \rightarrow \pi$ could in principle produce arbitrarily large amounts of DM, if isocurvature constraints can be avoided (at the cost of additional fine-tuning), as discussed in ref. [16].

Model E/N	DFSZAxion-I 8/3	DFSZAxion-II 2/3	KSVZAxion			
			0	2/3	5/3	8/3
Odds ($\Omega_a \lesssim \Omega_{\text{DM}}$)	3:1	1:1	1:2	1:2	1:2	1:2
Odds ($\Omega_a \sim \Omega_{\text{DM}}$)	1:5	1:6	1:6	1:6	1:6	1:6

Table 9. Odds ratios in favour of DFSZ- and KSVZ-type models, compared to the parent QCDAxion, as calculated from the nested sampling evidence estimates in MultiNest, and including the WD cooling hints. We either impose the DM relic density as an upper limit ($\Omega_a \lesssim \Omega_{\text{DM}}$) or demand that axions be all of DM ($\Omega_a \sim \Omega_{\text{DM}}$). Note that the estimated uncertainties on the evidence values are small enough that the corresponding uncertainties on the odds ratios are negligible.

parameter region in the DFSZAxion-I model is in fact larger, because the large values of C_{aee} needed to explain the cooling hints can be achieved more easily.

To investigate the consequences of fine-tuning in more detail, we consider all models in a Bayesian model comparison. The resulting odds ratios can be found in table 9. If we impose the DM relic density as an upper limit, the odds ratios are still mostly inconclusive. However, for the DFSZAxion-II and KSVZAxion models, the trend swings in favour of the broader class of QCDAxion models when cooling hints are added to the analysis. In contrast, the DFSZAxion-I model fares better than the QCDAxion, with an odds ratio of 3:1. If we combine this 3:1 odds ratio with the 2:1 preference for the QCDAxion model over KSVZAxion models, there is a 6:1 positive preference for the DFSZAxion-I model over all KSVZAxion models. This preference is not surprising, given the differences in the natural axion-electron coupling strength in these models.

If we demand that axions solve the cooling hints and constitute all of DM, model comparison confirms quantitatively (with ratios of $\leq 1:5$) that there is a positive preference for the QCDAxion over the DFSZ- and KSVZ-type models. This is because we allowed for much larger values of C_{aee} with QCDAxion models than with DFSZ- and KSVZ-type models. In fact, the one-dimensional marginal posterior for the QCDAxion electron coupling peaks at $C_{aee} \sim 100$ ($C_{aee} \sim 50$ if we impose the DM relic density as an upper limit), whereas the DFSZ and KSVZ models are limited to couplings of less than one. However, it should be noted that QCDAxion models with such a large coupling are not expected from traditional axion models, and may therefore pose a challenge for model building.

6 Conclusions

In this study we presented the first global fits of axion models in the pre-inflationary PQ symmetry-breaking scenario, using frequentist and Bayesian methods. We identified the most viable regions of parameter space for these models and discussed the effect of adding cooling hints seen in white dwarfs. We not only considered the phenomenological parameter space, but also the underlying parameters in a generic QCD axion model and six specific DFSZ- and KSVZ-type models. We extended previous results in the literature by including various nuisance parameters and by quantitatively considering the fine-tuning of the initial misalignment angle.

We found, in agreement with previous work, that QCD axion models are viable as explanations for both the observed cold dark matter and the white dwarf cooling hints. We showed that, for a broader class of QCDAxion models than what is traditionally considered, these can be achieved simultaneously. We also quantitatively confirmed that this is not possible for six specific DFSZ- and KSVZ-type models. However, if the condition of being all the dark matter in the Universe is relaxed, we found that DFSZAxion-I models are positively preferred over DFSZAxion-II models, and over all four KSVZAxion models that we investigated.

We determined the most credible predicted ranges for the QCD axion mass and its cosmological abundance in the Bayesian statistical framework. Although these results are somewhat prior-dependent, QCDAxions appear likely to be a cosmologically relevant (but probably not dominant) component of dark matter. Moreover, the most credible axion mass range is within reach of current and planned haloscope experiments.

Global fits of QCD axions and axion-like particles have the potential to confirm and refine previously known phenomenological statements about the relevant parameter spaces (e.g. exclusion limits, existence of fine-tuning, compatibility with cooling hints and dark matter). In other cases, they offer new and more rigorous insights (e.g. model comparison, most credible parameter regions). This is true in particular for the Bayesian analyses that we performed in this paper, because they inherently take fine-tuning into account. Due to the orthogonality of constraints and the insufficient sensitivity of most experiments, it is not (yet) possible to decisively probe the axion parameter space in the pre-inflationary PQ symmetry-breaking scenario, but it is realistically possible to target the most likely versions of QCD axion and axion-like particle models in the near future (see ref. [5] for a recent review of upcoming axion searches).

The axion routines that we developed for this paper will be publicly available in DarkBit within GAMBIT 1.3.0, to be released shortly. Statistical samples and input files from this study are freely available from Zenodo [35].

Acknowledgments

We are grateful to the U.K. Materials and Molecular Modelling Hub for computational resources, which is partially funded by EPSRC (EP/P020194/1), PRACE for awarding us access to Marconi at CINECA, Italy, and the Imperial College Research Computing Service as well as to all those individuals and experimental collaborations who kindly provided their data and results. We thank A. Serenelli for helpful comments on solar models, and acknowledge further useful discussions and communication with T. Battich, P. Brun, A. H. Córscico, I. Dominguez, M. Giannotti, I. G. Irastorza, J. Jaeckel, A. Lindner, L. di Luzio, D. J. E. Marsh, J. Redondo, A. Ringwald, O. Straniero, K. K. Szabo, J. Vogel, L. Rosenberg and G. Rybka. We also thank our colleagues within the GAMBIT community, in particular S. Bloor, G. Martinez and J. McKay. For all digitisation procedures, we acknowledge usage of the tool WebPlotDigitizer by A. Rohatgi, available at <https://automeris.io/WebPlotDigitizer/>. SH gratefully acknowledges funding by the Imperial College President's PhD Scholarship scheme. FK is supported by the DFG

Emmy Noether Grant No. KA 4662/1-1. PS is supported by STFC (ST/K00414X/1, ST/N000838/1, ST/P000762/1). MW is supported by the ARC Future Fellowship FT140100244.

A Integrating solar models for the signal prediction in CAST

To obtain the axion flux at Earth from axion-photon interactions, we need four inputs from a solar model: the solar radius $R_\odot$, the temperature $T(r)$, the plasma frequency $\omega_{\text{pl}}(r)$ and the screening scale $\kappa_s(r)$ (see eqs. 4.5–4.8). While $T(r)$ can be obtained directly from solar model files, $\omega_{\text{pl}}(r)$ and $\kappa_s(r)$ have to be calculated using the mass density $\rho(r)$ and the mass fractions $X_i(r)$ for each ion/atom with label i using (4.6). Assuming that the plasma is fully ionised, we can recast the equations into the following form [209]:

$$\kappa_s^2(r) = \frac{4\pi\alpha_{\text{EM}}}{T(r)} \frac{\rho(r)}{m_u} \sum_i \frac{X_i(r)Z_i}{A_i} (1 + Z_i) , \quad (\text{A.1})$$

$$\text{and } \omega_{\text{pl}}^2(r) = \frac{4\pi\alpha_{\text{EM}}}{m_e} \frac{\rho(r)}{m_u} \sum_i \frac{X_i(r)Z_i}{A_i} , \quad (\text{A.2})$$

where m_u is the atomic mass unit, r is the distance from the centre of the Sun in units of $R_\odot$, and Z_i and A_i are the charge and atomic weight of the i th element. Note that the approximation of full ionisation is justified almost everywhere inside the Sun, i.e. $r \lesssim 0.95$.²⁰ As the largest contribution to the axion flux comes from the innermost region ($r \lesssim 0.2$ [121]), we can safely employ this assumption. For elements tracked by the solar model in bulk (i.e. without isotopic information), we calculate the mean atomic weights A_i using the isotopic composition of ref. [161], with values based on the terrestrial composition updated to use more recent data [232].

The expected number of photons in the energy range $[E_j, E_{j+1}]$ is

$$s_j = \int_{E_j}^{E_{j+1}} \mathcal{E}(E) \frac{d\Phi(E)}{dE} dE , \quad (\text{A.3})$$

where $\mathcal{E}$ is the effective exposure and $d\Phi/dE$ is the combined photon spectrum from axion-electron and axion-photon contributions. We neglect the energy dispersion of the CAST detector, as it is always less than about 0.2 keV [121], and therefore smaller than or comparable to the bin width of the CAST analyses (0.3 keV and 0.5 keV respectively for the 2007 and 2017 analyses). We provide tabulated data within DarkBit for the effective exposure $\mathcal{E}$;²¹ this can be found in DarkBit/data as `datasetEffectiveExposure.dat`, where `dataset` is either CAST2007 or CAST2017_`X`, where `X` corresponds to the data sets A to L in ref. [122]. All other files that we mention in the following can be found in the same folder.

Note that performing the energy integration in (A.3) after the density integral in (4.8) is only possible because all contributions to the integral are for energies greater than $\omega_{\text{pl}}(0) \approx 0.3$ keV. For energies lower than ω_{pl} , axions cannot be produced from the plasma and the square root in (4.8) becomes ill-defined.

²⁰Aldo Serenelli, private communication.

²¹I. Irastorza and J. Vogel, private communication.

To obtain the contribution to $d\Phi/dE$ from axion-photon interactions, we integrate (4.7) and (4.8) over r and ρ , using the adaptive 51 point Gauss-Kronrod rule `gsl_integration_qag`, from the `gsl` library. We obtain quantities necessary for these integrations (temperature, etc.) by interpolating the solar model linearly in radius. For the contribution to $d\Phi/dE$ from axion-electron interactions, we use the spectrum published in ref. [146] and redistribute it as `Axion_Spectrum_AGSS09met_old_gaee.dat`.

The peaks in the spectrum from axion-electron interactions specifically require using an algorithm that takes them into account as singularities. We therefore compute the contribution from axion-electron interactions to the signals s_j (A.3) using the `gsl_integration_qagg` integrator of the `gsl` library.²² For the axion-photon contribution to the signal, we again use the 51-point Gauss-Kronrod rule. We perform all signal calculations to a relative accuracy of 10^{-6} .

We compute signals at reference values of $g_{a\gamma\gamma} = 10^{-10} \text{ GeV}^{-1}$ and $g_{aee} = 10^{-13}$, at 183 mass values ranging from 10^{-3} to 10^2 eV . We use unequally spaced mass values, as the density of points needs to be higher in certain regions to obtain good interpolating functions. We provide these pre-calculated signal count files as `dataset.ReferenceCounts_solarmodel_coupling.dat`, where `dataset` refers to CAST2007 or CAST2017_X results, `solarmodel` is the name of the solar model, and `coupling` is either the contribution from axion-photon interactions (`gagg`) or axion-electron interactions (`gaee`).

For the axion-photon contribution (`coupling = gagg`), we provide these files for the GS98 [162, 163], AGSS09ph [160], AGSS09met_old [160] and AGSS09met [162] models. For the axion-electron contribution (`coupling = gaee`), where a direct calculation of $d\Phi/dE$ is not easily performed, we provide a single file based on the earlier iteration of the AGSS09met model (`AGSS09met_old`); this corresponds to the model used in ref. [146] to compute the spectrum that we ship in DarkBit.

The user may choose which solar model to employ from the `Rules` section of the YAML file, using the keys `solar_model_gagg` and `solar_model_gaee`, with `dataset` being CAST2007 or CAST2017 (no suffix in this case, as the CAST 2017 likelihood combines all the results):

```
Rules:
- capability: dataset_signal_vac
  options:
    solar_model_gagg: solarmodel
    solar_model_gaee: solarmodel
```

If these options are left unspecified, the defaults will be chosen. For the axion-photon coupling, this is AGSS09met [162]; for the axion-electron coupling, the default is AGSS09met_old [160]. If the corresponding `dataset.ReferenceCounts_solarmodel_coupling.dat` file cannot be found, DarkBit will attempt to create such a file by solving (A.3). For the axion-photon coupling, it will attempt to do this on the basis of a provided solar model file, which should be called `SolarModel_solarmodel.dat`. We have included the solar model file `SolarModel_AGSS09met.dat` from ref. [162] as an example of the expected formatting. For the axion-electron coupling, DarkBit will attempt to generate the relevant

²²Note that the `gsl` library is already required by GAMBIT.

`dataset.ReferenceCounts_solarmodel_gaee.dat` file on the basis of a provided spectrum file `Axion.Spectrum_solarmodel_gaee.dat`. If the relevant solar model or spectrum file cannot be located, or has the wrong format, GAMBIT will terminate. Note that this will currently happen for any choice of `solar_model_gaee` except for the default value. However, if the user can provide the spectrum file for any other model, the calculation will go ahead.

B Numerical implementation of the solution to the axion field equation

The most general equation of motion for a scalar field $\phi(t, \mathbf{x})$ in the background of a Friedmann-Robertson-Walker-Lemaître universe with Hubble parameter $H \equiv \dot{a}/a$ and potential V is given by [59, appendix B.12]

$$\ddot{\phi} + 3H(t)\dot{\phi} + \frac{\Delta\phi}{a^2(t)} + V'[\phi] = 0, \quad (\text{B.1})$$

while the energy density associated with ϕ can be calculated using:

$$\varepsilon = \frac{1}{2}\dot{\phi}^2 - \frac{1}{2a^2(t)}(\nabla\phi)^2 + V[\phi]. \quad (\text{B.2})$$

To obtain the energy density in axions today, it is sufficient to solve the *homogeneous* field equation for $\theta \equiv a/f_a$, the normalised axion field:

$$\ddot{\theta} + 3H\dot{\theta} + \frac{V'[\theta]}{f_a^2} = 0. \quad (\text{B.3})$$

We assume the potential

$$V[\theta] = f_a^2 m_a^2 [1 - \cos(\theta)], \quad (\text{B.4})$$

for QCD axions and ALPs alike. For QCD axions, this is a reasonable approximation to the full potential [51]. The axion energy density $\varepsilon_a \simeq \rho_a$ from (B.2) is therefore given by

$$\rho_a = \frac{1}{2}f_a^2 \dot{\theta}^2 + f_a^2 m_a^2 [1 - \cos(\theta)]. \quad (\text{B.5})$$

The initial conditions for (B.3) are some initial misalignment angle $\theta(0) \equiv \theta_i$ and vanishing initial derivatives $\dot{\theta}(0) \equiv 0$.²³ For early times, the solution of (B.3) is a constant value, $\theta(t) = \theta_i$. The behaviour only changes around the oscillation time T_{osc} , defined as the point when $3H(T_{\text{osc}}) = m_a(T_{\text{osc}})$.²⁴

Assuming that the scale factor a is a monotonically increasing function of cosmic time t , and that the temperature T is a monotonically decreasing function of t , we may re-write (B.3) as

$$D_X\theta + \frac{V'[\theta]}{f_a^2} = 0, \quad (\text{B.6})$$

²³The initial value for the derivative can be taken as zero assuming that $|\dot{\theta}/H| \ll 1$ at early times [68, ch. 10.3.2]. The consequence of that assumption is essentially independent of the cosmological history and equation of state [179].

²⁴We will solve the field equation numerically around this point and the exact definition of the onset of oscillation is therefore irrelevant. See e.g. ref. [40] for a more detailed discussion of the general relevance of this definition.

where $X = t, \alpha, \tau$ with $\alpha \equiv a/a_{\text{osc}}, \tau \equiv T/T_{\text{osc}}$ and

$$D_t = \frac{d^2}{dt^2} + 3H \frac{d}{dt}, \quad (\text{B.7})$$

$$D_\alpha = H^2 \alpha^2 \left[\frac{d^2}{d\alpha^2} + \left(\frac{1}{\alpha} + \frac{1}{H} \frac{dH}{d\alpha} \right) \frac{d}{d\alpha} \right], \quad (\text{B.8})$$

$$D_\tau = \left(\frac{H\alpha}{\alpha^\times} \right)^2 \left[\frac{d^2}{d\tau^2} + \left(\frac{H^\times}{H} + \frac{\alpha^\times}{\alpha} - \frac{\alpha^{\times\times}}{\alpha^\times} \right) \frac{d}{d\tau} \right], \quad (\text{B.9})$$

where $^\times$ denotes derivatives w.r.t. τ . Conservation of entropy can subsequently be used to relate α and temperature:

$$\alpha = \left(\frac{g_S(T)}{g_S(T_{\text{osc}})} \right)^{-1/3} \tau^{-1} \equiv \gamma(\tau)/\tau. \quad (\text{B.10})$$

Equation (B.9) has the advantage that an explicitly temperature-dependent axion mass or the effective relativistic degrees of freedom, g and g_S , can be easily incorporated. The various terms can be further simplified as follows:

$$\frac{\alpha^\times}{\alpha} = \frac{\gamma^\times}{\gamma} - \frac{1}{\tau}, \quad (\text{B.11})$$

$$\frac{\alpha^{\times\times}}{\alpha^\times} = \frac{\alpha^{\times\times}}{\alpha} \frac{\alpha}{\alpha^\times} = \frac{\tau \gamma^{\times\times}}{\tau \gamma^\times - \gamma} - \frac{2}{\tau}. \quad (\text{B.12})$$

Rescaling $\vartheta \equiv \theta/\theta_i$, we may use (B.9) to rewrite (B.6) as

$$\vartheta^{\times\times} + \frac{F_2(\tau)}{\tau} \vartheta^\times + \left(\frac{F_1(\tau) m_a(\tau)}{\tau H(\tau)} \right)^2 \frac{\sin(\theta_i \vartheta)}{\theta_i} = 0, \quad (\text{B.13})$$

where the auxiliary functions F_i for $i = 1, 2$ are given by

$$F_1(\tau) = -1 + \frac{\tau \gamma^\times}{\gamma} = - \left(1 + \frac{\tau g_S^\times}{3 g_S} \right), \quad (\text{B.14})$$

$$F_2(\tau) = 2 + \frac{\tau H^\times}{H} - \frac{\tau^2 g_S^{\times\times} + 4\tau g_S^\times}{\tau g_S^\times + 3g_S}, \quad (\text{B.15})$$

$$= \frac{\tau g^\times}{2 g} - \frac{\tau^2 g_S^{\times\times} + 4\tau g_S^\times}{\tau g_S^\times + 3g_S}, \quad (\text{B.16})$$

where we assume a radiation-dominated universe, $H \propto \sqrt{g(T)} T^2$, for the last line. We use interpolated values for (B.14) and (B.16) based on the analytic forms for the effective degrees of freedom in ref. [64]. The tables used for this procedure are included in **GAMBIT**.

Note that if the changes in the effective degrees of freedom are not important, $F_1(\tau) \approx -1$, $F_2(\tau) \approx 0$, and the field equation (B.13) reads in the harmonic limit of $|\theta_i| \ll 1$:

$$\theta^{\times\times} \approx - \left(\frac{m_a(\tau)}{H\tau} \right)^2 \theta, \quad (\text{B.17})$$

i.e. we obtain the approximate behaviour of a (damped) harmonic oscillator. In that limit, the comoving axion number density $n_a \equiv \rho_a a^3/m_a$ is conserved on average. We can

therefore stop integrating the differential equation once specified conditions for $m_a/3H$ and θ are met. Because the solution $\theta(t)$ oscillates as a function of time or temperature, we define a peak amplitude $\hat{\theta}$ in order to check that $|\theta(t_*)| \ll 1$ from some time t_* :

$$\hat{\theta}(t) \equiv \frac{\sqrt{2\rho_a(t)}}{f_a m_a(t)}. \quad (\text{B.18})$$

In the harmonic limits, (B.18) indeed coincides with the amplitude of oscillation when all the energy of the field is stored in the potential part of the energy density.²⁵ Once the axion field starts to oscillate, fulfilment of the condition $\hat{\theta} \ll 1$ indicates that the harmonic limit has been achieved. This provides much clearer limiting behaviour than e.g. $|\theta|$.

We use the `gsl` implementation of Brent's method to solve for T_{osc} to a precision of 10^{-6} . We then impose the initial conditions at some relatively high temperature $\tau_{\text{start}} = 10^5$, and evolve the differential equation with `gsl_odeiv2_step_bsimp` until some relative temperature τ_{stop} for which $m_a/3H > 10^3$ and $\hat{\theta} < 10^{-2}$. We then calculate the energy density from realignment axions today using the conservation of comoving number density

$$\rho_a = \frac{m_{a,0}}{m_{a,*}} \frac{g_S(T_{\text{CMB}})}{g_S(T_*)} \left(\frac{T_{\text{CMB}}}{T_*} \right)^3 \rho_{a,*}, \quad (\text{B.19})$$

where T_{CMB} is the CMB temperature.

C H.E.S.S. likelihood implementation details

Our approximation of the H.E.S.S. likelihood is based on the information provided in figures 6 and 7 of ref. [133] on limits on axion-photon conversion in galactic cluster magnetic fields. We choose 13 lines of constant axion-photon coupling (corresponding to specific values of the axion decay rate Γ in their figure) that intersect with the exclusion curves at least five times (dashed black and blue horizontal lines and blue circles in figure 28).

We subsequently define a family of one-dimensional natural cubic splines along each dashed line by including two additional points *outside* of the known data. We optimise each of those splines by iteratively modifying the locations of the additional points along the dashed lines, until the values and first derivatives of the splines at the additional points go to zero (blue diamonds in figure 28). This forcibly prevents ringing, ensuring that the interpolated log-likelihood is negative everywhere. To evaluate the log-likelihood at a given axion mass on one of the horizontal dashed lines, we use the resultant interpolating cubic spline from the fitting procedure. If a coupling value of interest sits between two horizontal dashed lines, we interpolate linearly between the values at a given mass on the adjacent dashed lines. We assign zero log-likelihood to all points outside of the area defined by the exterior points (blue diamonds).

This leaves two more regions in figure 28. The first is at higher couplings, where figure 6 of ref. [133] does not show any limits, but figure 7 does (dotted horizontal lines

²⁵Note that due to the inequality $1 - \cos(x) \leq x^2/2$, we can apply this procedure even if θ is not much smaller than unity. In fact, it should be valid even if the true axion potential is not cosine-shaped, as long as the true potential is also harmonic for small values of θ .

Figure 28. Detailed description of our implementation of the H.E.S.S. exclusion limits from gamma-ray spectral distortions of PKS 2155-304 (figure 8). We show the points that form the basis of our interpolation: points directly from the exclusion curves (all blue triangles, circles and stars), estimated points (squares) as well as points directly inferred from the interpolation scheme (diamonds). See the main text for further details.

and blue triangles in figure 28). To complete the likelihood curve families in this region, we infer additional points (by eye; light blue squares in figure 28) and follow the procedure described in the previous paragraph. We assign zero log-likelihood to all couplings above the uppermost horizontal line in figure 28.

At low couplings, less than five contours are available for interpolating (inverted triangles in figure 28). Here, we construct a single cubic spline in the vertical direction (blue stars and vertical dot-dash line in figure 28), in order to determine an exterior point at low coupling (the diamond at the bottom of the plot). We define three parabolae, based on nine points: the set of two diamonds and four outer circles on the dashed black line around $g_{a\gamma\gamma} = 10^{-10.6} \text{ GeV}^{-1}$, the two stars below this line, and the one diamond at the bottom of the plot. We designate the lower parabola, passing through the exterior point (bottommost diamond) at low coupling and drawn as a dashed black curve, as the zero-log-likelihood contour. We use the middle parabola to model the 99% C.L. exclusion curve, and the upper one to model the 95% C.L. curve. These three parabolae allow us to map any point (represented by the white star) inside the lower parabola to a likelihood, by interpolating between the three parabolae. We do this using the one-dimensional spline constructed along the central vertical axis, using as input to the interpolation the position of the point in question between the central star and the zero-log-likelihood parabola. We carry out all interpolations in log space for both the axion mass and the axion-photon coupling.

D Prior-dependence of results

For QCD axions some prior choices are relatively straightforward (such as the use of a flat prior for θ_i) or do not have a significant impact on the results (such as the prior

Figure 29. Marginalised posteriors for QCDaxion models with three alternative priors. Here we impose the observed DM relic density as an upper limit. We show the one-dimensional posterior for the axion mass without the cooling hints (*left*; to be compared to the right panel of figure 17) and the two-dimensional posterior for $m_{a,0}$ vs $g_{a\gamma\gamma}$, including the cooling hints (*right*; to be compared with figure 25). The priors adopted in the main part of this study are marked in grey in the left panel, and in the colour scale and black contours in the right panel.

assignments for the well-constrained nuisance parameters). In general, however, the prior beliefs can have a significant impact on the posteriors and Bayes factors, especially in the absence of constraining data. In this appendix, we investigate different prior choices for the parameters f_a , C_{aee} , and E/N in order to assess the robustness of our results and to understand how our conclusions may be affected by a variation of priors. All the new results that we present in this appendix are based on **MultiNest** runs only, as we need the evidence values for examining the impacts of priors on model comparison, but are not sufficiently interested in the finer details of posterior maps to warrant also running **T-Walk**.

First, let us revisit the determination of the most credible regions for the axion mass in QCDaxion models from section 5.3. The left panel of figure 29 shows the marginalised posterior distribution for the axion mass, $m_{a,0}$, for a number of alternative priors. The grey line and shading correspond to the prior choice made in the main text (right panel of figure 17). We can see that choosing a flat prior on either C_{aee} or f_a shifts the most credible region for the axion mass to lower masses. The effect is more extreme for a flat prior on f_a than for a flat prior on C_{aee} . On the other hand, restricting the prior of the anomaly ratio to $E/N \leq 170/3$ (the highest values achievable in KSVZ-type models) does not have a significant impact on the result.

The different prior choices also alter the most credible regions of the axion-photon couplings (cf. figure 25). The resulting posterior distributions are shown in the right panel of figure 29. As expected, priors that imply a preference for smaller axion masses also lead to a preference for smaller values of $g_{a\gamma\gamma}$, while changes in prior that do not significantly affect the most probable region for $m_{a,0}$ also do not impact the findings for $g_{a\gamma\gamma}$ very much. In particular, for the alternative priors on E/N and C_{aee} , the regions of highest posterior density still overlap with the region that we obtained with our adopted prior, whereas the regions resulting from a flat prior on f_a are almost completely disjoint from the others.

The preference for very small axion masses for the case of a flat prior on f_a is a direct consequence of the relation between f_a and $m_{a,0}$, which implies that the prior probability for $m_{a,0}$ is strongly peaked at the lowest possible axion masses. This prior then overwhelms the preference for large $m_{a,0}$ from the combination of tuning in the initial misalignment angle and the relic density requirement. In other words, the preferred mass range in this case is dictated almost entirely by the assumed prior. A log prior on f_a , on the other hand, essentially corresponds to a log prior on $m_{a,0}$, which has the important advantage that any preference in the axion mass range is the result of phenomenological requirements on the model. A logarithmic prior for f_a is therefore less informative than a flat one, which favours a particular scale of new physics near the upper end of the prior range.

However, it is also important to determine how changing the range of the logarithmic prior impacts our results. Taking the example of the grey curve in the left panel of figure 29 (canonical priors, no cooling hints, $\Omega_a \lesssim \Omega_{\text{DM}}$), we find that restricting the log prior to $f_a \in [10^7, 10^{15}]$ GeV results in a 95% equal-tailed credible interval for the axion mass of $0.50 \mu\text{eV} \leq m_{a,0} \leq 3.8 \text{ meV}$, compared to $0.48 \mu\text{eV} \leq m_{a,0} \leq 3.8 \text{ meV}$ for the prior range that we adopted in the main paper ($f_a \in [10^6, 10^{16}]$ GeV). Narrowing this range further to $f_a \in [10^8, 10^{14}]$ GeV only raises the lower edge of the interval slightly: $0.63 \mu\text{eV} \leq m_{a,0} \leq 3.6 \text{ meV}$. Clearly the range of the logarithmic prior on f_a has little impact on our results. The median values and credible intervals of Ω_a are even more stable, with the lower bound of the 95% equal-tailed credible interval slightly increasing from 6.8×10^{-6} to 7.3×10^{-6} for $f_a \in [10^8, 10^{14}]$ GeV (corresponding to about 0.006% of the cosmological density of DM, regardless of the prior range).

We also calculate the Bayes factors for KSVZ- and DFSZ-type models compared to QCDAxion models, adopting the alternative priors for the appropriate models. The resulting odds ratios are given in table 10, in addition to the values for the priors that we adopt in section 5.

As in figure 29, adopting a reduced range for the anomaly ratio E/N has little impact, slightly increasing the odds in favour of KSVZ-type models. Choosing the alternative prior on the axion-electron coupling constant, C_{aee} , the odds in favour of KSVZ- and DFSZ-type models increase by an order of magnitude (if cooling hints are not included). This is because the flat prior on C_{aee} causes a preference for large f_a (smaller $m_{a,0}$). This, in turn, implies a considerable fine-tuning in θ_i in order to avoid overproducing DM. After including the cooling hints, axion-electron interactions are much more tightly constrained by data, and the prior-dependence of the odds ratios is reduced.

Finally, a flat prior on f_a favours the lowest allowed axion masses in all models. This results in even more fine-tuning of the initial misalignment angle than using a flat prior on C_{aee} for QCDAxion models. However, although the Bayesian evidence in QCDAxion models is drastically reduced, the odds ratios in favour of KSVZ- and DFSZ-type models turn out to be 1:1. This is because the alternative prior is applied to *all* models: the fine-tuning required to overcome the prior pushes all evidences equally low, washing out any preference one way or another from the data.

We see that prior choice can indeed significantly influence the results of our analysis. This is true in particular for a flat prior on f_a , where the prior-dependence dominates the

Model	DFSZAxion-I	DFSZAxion-II	KSVZAxion			
E/N	8/3	2/3	0	2/3	5/3	8/3
Priors adopted in this study						
Odds (w/o WD cooling)	2:1	2:1	2:1	2:1	2:1	2:1
Odds (with WD cooling)	3:1	1:1	1:6	1:6	1:6	1:6
Reduced prior range on E/N in QCDAxion models						
Odds (w/o WD cooling)	1:1	2:1	2:1	2:1	2:1	2:1
Odds (with WD cooling)	2:1	1:2	1:4	1:4	1:3	1:4
Flat prior on C_{aee} in QCDAxion models						
Odds (w/o WD cooling)	23:1	27:1	25:1	26:1	31:1	28:1
Odds (with WD cooling)	3:1	1:1	1:2	1:2	1:2	1:2
Flat prior on f_a in <i>all</i> models						
Odds (w/o WD cooling)	1:1	1:1	1:1	1:1	1:1	1:1
Odds (with WD cooling)	1:1	1:1	1:1	1:1	1:1	1:1

Table 10. Odds ratios in favour of DFSZ- and KSVZ-type models, compared to the parent QCDAxion model, as calculated from nested sampling evidence estimates by MultiNest. We impose the DM relic density as an upper limit ($\Omega_a \lesssim \Omega_{\text{DM}}$). Note that the estimated uncertainties on the evidence values are small enough that the corresponding uncertainties on the odds ratios are negligible.

results. On the other hand, we also saw that data with a strong preference for certain model parameter values, such as the cooling hints, can reduce the prior-dependence. In the face of more such data, the impact of prior choices tends to become less pronounced, ideally leaving only the physically-motivated prior on the initial misalignment angle able to influence final results.

E Overview of new capabilities

For reference, in table 11 we provide a complete list of the new capabilities, dependencies and options that we have added to DarkBit whilst preparing this paper.

Capability	Function (Return Type):	Dependencies
	Brief Description	
QCDAxion_ZeroTemperatureMass	QCDAxion_ZeroTemperatureMass_Nuisance_lnl (double) : Nuisance likelihood for the QCD-related energy scale Λ_χ , which defines the low-energy QCD axion mass.	
QCDAxion_TemperatureDependence	QCDAxion_TemperatureDependence_Nuisance_lnl (double) : Nuisance likelihood for the parameters β and T_χ , which define the temperature dependence of the QCD axion mass for high temperatures.	
QCDAxion_AxionPhotonConstant	QCDAxion_AxionPhotonConstant_Nuisance_lnl (double) : Nuisance likelihood for the contribution of axion-pion interactions in the chiral Lagrangian to the axion-photon coupling.	
ALPS1_signal_vac	calc_ALPS1_signal_vac (double) : Calculates the signal prediction for one data frame of the ALPS 1 experiment (evacuated setup).	
ALPS1_signal_gas	calc_ALPS1_signal_gas (double) : Calculates the signal prediction for one data frame of the ALPS 1 experiment (setup filled with Argon gas).	ALPS1_signal_vac, ALPS1_signal_gas
lnL_ALPS1	calc_lnL_ALPS1 (double) : Log-likelihood for all runs of the ALPS 1 experiment.	
CAST2007_signal_vac*	calc_CAST2007_signal_vac (std::vector<double>) : Signal prediction for a 2004 vacuum run of the CAST experiment.	
lnL_CAST2007	calc_lnL_CAST2007 (double) : Log-likelihood for CAST 2007 experimental results.	CAST2007_signal_vac
CAST2017_signal_vac*	calc_CAST2017_signal_vac (std::vector<double>) : Signal prediction for the combined 2013, 2014 and 2015 vacuum runs of the CAST experiment.	
lnL_CAST2017	calc_lnL_CAST2017 (double) : Log-likelihood for CAST 2017 experimental results.	CAST2017_signal_vac
Haloscope_signal	calc_Haloscope_signal (double) : Calculates the relative expected signal for a haloscope ignoring experimental parameters and mass dependence.	
lnL_Haloscope_ADMX1	calc_lnL_Haloscope_ADMX1 (double) : Log-likelihood for the combined ADMX 1998 to 2011 experimental results.	Haloscope_signal
lnL_Haloscope_ADMX2	calc_lnL_Haloscope_ADMX2 (double) : Log-likelihood for the ADMX 2018 experimental result.	Haloscope_signal
lnL_Haloscope_RBF	calc_lnL_Haloscope_RBF (double) : Log-likelihood for the RBF experiment.	Haloscope_signal
lnL_Haloscope_UF	calc_lnL_Haloscope_UF (double) : Log-likelihood for the UF experiment.	Haloscope_signal
AxionOscillationTemperature	calc_AxionOscillationTemperature (double) : Calculates the temperature scale at which the axion field starts to oscillate.	
RD_oh2†	RD_oh2_Axions (double) : The axion dark matter relic density.	
RParameter	calc_RParameter (double) : Calculates the expected R parameter for a system with axions.	
lnL_RParameter	calc_lnL_RParameter (double) : Log-likelihood for the observed R parameter.	RParameter
lnL_HESS_GCMF	calc_lnL_HESS_GCMF (double) : Log-likelihood for H.E.S.S. constraints on axions (galactic cluster magnetic field result).	
lnL_SN1987A	calc_Haloscope_signal (double) : Log-likelihood for an axion signal from SN1987A.	
lnL_WDVar_G117B15A	calc_lnL_WDVar_G117B15A (double) : Log-likelihood for the white dwarf G117-B15A (cooling hint).	
lnL_WDVar_R548	calc_lnL_WDVar_R548 (double) : Log-likelihood for the white dwarf R548 (cooling hint).	
lnL_WDVar_PG1351489	lnL_WDVar_PG1351489 (double) : Log-likelihood for the white dwarf PG 1351+489 (cooling hint).	
lnL_WDVar_L192	lnL_WDVar_L192 (double) : Log-likelihood for the white dwarf L192 (cooling hint).	

* Options explained in Appendix ??

† New module function able to deliver a pre-existing capability

Table 11. New, axion-related capabilities for the DarkBit module in GAMBIT.

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

References

- [1] R.D. Peccei and H.R. Quinn, *CP conservation in the presence of instantons*, *Phys. Rev. Lett.* **38** (1977) 1440 [[INSPIRE](#)].
- [2] R.D. Peccei and H.R. Quinn, *Constraints imposed by CP conservation in the presence of instantons*, *Phys. Rev. D* **16** (1977) 1791 [[INSPIRE](#)].
- [3] S. Weinberg, *A new light boson?*, *Phys. Rev. Lett.* **40** (1978) 223 [[INSPIRE](#)].
- [4] F. Wilczek, *Problem of strong P and T invariance in the presence of instantons*, *Phys. Rev. Lett.* **40** (1978) 279 [[INSPIRE](#)].
- [5] I.G. Irastorza and J. Redondo, *New experimental approaches in the search for axion-like particles*, *Prog. Part. Nucl. Phys.* **102** (2018) 89 [[arXiv:1801.08127](#)] [[INSPIRE](#)].
- [6] J. Preskill, M.B. Wise and F. Wilczek, *Cosmology of the invisible axion*, *Phys. Lett. B* **120** (1983) 127 [[INSPIRE](#)].
- [7] L.F. Abbott and P. Sikivie, *A cosmological bound on the invisible axion*, *Phys. Lett. B* **120** (1983) 133 [[INSPIRE](#)].
- [8] M. Dine and W. Fischler, *The not so harmless axion*, *Phys. Lett. B* **120** (1983) 137 [[INSPIRE](#)].
- [9] M.S. Turner, *Cosmic and local mass density of “invisible” axions*, *Phys. Rev. D* **33** (1986) 889.
- [10] J. Isern, M. Hernanz and E. Garcia-Berro, *Axion cooling of white dwarfs*, *Astrophys. J.* **392** (1992) L23 [[INSPIRE](#)].
- [11] A.H. Corsico 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 [[arXiv:1205.6180](#)] [[INSPIRE](#)].
- [12] A.H. Corsico et al., *An independent limit on the axion mass from the variable white dwarf star R548*, *JCAP* **12** (2012) 010 [[arXiv:1211.3389](#)] [[INSPIRE](#)].
- [13] M. Giannotti, I. Irastorza, J. Redondo and A. Ringwald, *Cool WISPs for stellar cooling excesses*, *JCAP* **05** (2016) 057 [[arXiv:1512.08108](#)] [[INSPIRE](#)].
- [14] A.H. Córscico 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* **07** (2016) 036 [[arXiv:1605.06458](#)] [[INSPIRE](#)].
- [15] T. Battich et al., *First axion bounds from a pulsating helium-rich white dwarf star*, *JCAP* **08** (2016) 062 [[arXiv:1605.07668](#)] [[INSPIRE](#)].
- [16] M. Giannotti et al., *Stellar recipes for axion hunters*, *JCAP* **10** (2017) 010 [[arXiv:1708.02111](#)] [[INSPIRE](#)].
- [17] 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 [[arXiv:0707.4312](#)] [[INSPIRE](#)].
- [18] M. Simet, D. Hooper and P.D. Serpico, *The Milky Way as a kiloparsec-scale axionscope*, *Phys. Rev. D* **77** (2008) 063001 [[arXiv:0712.2825](#)] [[INSPIRE](#)].

- [19] I.F.M. Albuquerque and A. Chou, *A faraway quasar in the direction of the highest energy auger event*, *JCAP* **08** (2010) 016 [[arXiv:1001.0972](#)] [[INSPIRE](#)].
- [20] A. De Angelis, G. Galanti and M. Roncadelli, *Relevance of axion-like particles for very-high-energy astrophysics*, *Phys. Rev. D* **84** (2011) 105030 [Erratum *ibid.* **D 87** (2013) 109903] [[arXiv:1106.1132](#)] [[INSPIRE](#)].
- [21] D. Horns and M. Meyer, *Indications for a pair-production anomaly from the propagation of VHE gamma-rays*, *JCAP* **02** (2012) 033 [[arXiv:1201.4711](#)] [[INSPIRE](#)].
- [22] 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 [[arXiv:1302.1208](#)] [[INSPIRE](#)].
- [23] GAMBIT collaboration, *GAMBIT: the Global And Modular Beyond-the-standard-model Inference Tool*, *Eur. Phys. J. C* **77** (2017) 784 [[arXiv:1705.07908](#)] [[INSPIRE](#)].
- [24] GAMBIT collaboration, *ColliderBit: a GAMBIT module for the calculation of high-energy collider observables and likelihoods*, *Eur. Phys. J. C* **77** (2017) 795 [[arXiv:1705.07919](#)] [[INSPIRE](#)].
- [25] THE GAMBIT DARK MATTER WORKGROUP collaboration, *DarkBit: a GAMBIT module for computing dark matter observables and likelihoods*, *Eur. Phys. J. C* **77** (2017) 831 [[arXiv:1705.07920](#)] [[INSPIRE](#)].
- [26] THE GAMBIT FLAVOUR WORKGROUP collaboration, *FlavBit: a GAMBIT module for computing flavour observables and likelihoods*, *Eur. Phys. J. C* **77** (2017) 786 [[arXiv:1705.07933](#)] [[INSPIRE](#)].
- [27] GAMBIT MODELS WORKGROUP collaboration, *SpecBit, DecayBit and PrecisionBit: GAMBIT modules for computing mass spectra, particle decay rates and precision observables*, *Eur. Phys. J. C* **78** (2018) 22 [[arXiv:1705.07936](#)] [[INSPIRE](#)].
- [28] GAMBIT collaboration, *Comparison of statistical sampling methods with ScannerBit, the GAMBIT scanning module*, *Eur. Phys. J. C* **77** (2017) 761 [[arXiv:1705.07959](#)] [[INSPIRE](#)].
- [29] GAMBIT collaboration, *A global fit of the MSSM with GAMBIT*, *Eur. Phys. J. C* **77** (2017) 879 [[arXiv:1705.07917](#)] [[INSPIRE](#)].
- [30] GAMBIT collaboration, *Status of the scalar singlet dark matter model*, *Eur. Phys. J. C* **77** (2017) 568 [[arXiv:1705.07931](#)] [[INSPIRE](#)].
- [31] GAMBIT collaboration, *Global fits of GUT-scale SUSY models with GAMBIT*, *Eur. Phys. J. C* **77** (2017) 824 [[arXiv:1705.07935](#)] [[INSPIRE](#)].
- [32] P. Athron et al., *Impact of vacuum stability, perturbativity and XENON1T on global fits of $\mathbb{Z}_2$ and $\mathbb{Z}_3$ scalar singlet dark matter*, *Eur. Phys. J. C* **78** (2018) 830 [[arXiv:1806.11281](#)] [[INSPIRE](#)].
- [33] GAMBIT collaboration, *Global analyses of Higgs portal singlet dark matter models using GAMBIT*, *Eur. Phys. J. C* **79** (2019) 38 [[arXiv:1808.10465](#)] [[INSPIRE](#)].
- [34] GAMBIT collaboration, *Combined collider constraints on neutralinos and charginos*, [arXiv:1809.02097](#) [[INSPIRE](#)].
- [35] GAMIT collaboration, *Supplementary Data: axion global fits with Peccei-Quinn symmetry breaking before inflation using GAMBIT*, available on [zenodo](#) (2018).
- [36] J.E. Kim, *Light pseudoscalars, particle physics and cosmology*, *Phys. Rept.* **150** (1987) 1.
- [37] M. Kuster, G. Raffelt and B. Beltran, *Axions: theory, cosmology, and experimental searches. Proceedings, 1st Joint ILIAS-CERN-CAST axion training, Geneva, Switzerland, November 30-December 2, 2005, Lect. Notes Phys.* **741** (2008) 1.

- [38] P. Arias et al., *WISPy cold dark matter*, *JCAP* **06** (2012) 013 [[arXiv:1201.5902](#)] [[INSPIRE](#)].
- [39] M. Kawasaki and K. Nakayama, *Axions: theory and cosmological role*, *Ann. Rev. Nucl. Part. Sci.* **63** (2013) 69 [[arXiv:1301.1123](#)] [[INSPIRE](#)].
- [40] D.J.E. Marsh, *Axion cosmology*, *Phys. Rept.* **643** (2016) 1 [[arXiv:1510.07633](#)] [[INSPIRE](#)].
- [41] J.E. Kim, *Weak interaction singlet and strong CP invariance*, *Phys. Rev. Lett.* **43** (1979) 103 [[INSPIRE](#)].
- [42] M.A. Shifman, A.I. Vainshtein and V.I. Zakharov, *Can confinement ensure natural CP invariance of strong interactions?*, *Nucl. Phys. B* **166** (1980) 493 [[INSPIRE](#)].
- [43] A.R. Zhitnitskij, *On possible suppression of the axion-hadron interactions*, *Sov. J. Nucl. Phys.* **31** (1980) 260.
- [44] M. Dine, W. Fischler and M. Srednicki, *A simple solution to the strong CP problem with a harmless axion*, *Phys. Lett.* **104B** (1981) 199 [[INSPIRE](#)].
- [45] M.D. Schwartz, *Quantum field theory and the standard model*, Cambridge University Press, Cambridge U.K. (2014).
- [46] R.J. Crewther, P. Di Vecchia, G. Veneziano and E. Witten, *Chiral estimate of the electric dipole moment of the neutron in quantum chromodynamics*, *Phys. Lett.* **88B** (1979) 123 [*Erratum ibid.* **B 91** (1980) 487] [[INSPIRE](#)].
- [47] M. Pospelov and A. Ritz, *Theta induced electric dipole moment of the neutron via QCD sum rules*, *Phys. Rev. Lett.* **83** (1999) 2526 [[hep-ph/9904483](#)] [[INSPIRE](#)].
- [48] J.M. Pendlebury et al., *Revised experimental upper limit on the electric dipole moment of the neutron*, *Phys. Rev. D* **92** (2015) 092003 [[arXiv:1509.04411](#)] [[INSPIRE](#)].
- [49] C. Vafa and E. Witten, *Restrictions on symmetry breaking in vector-like gauge theories*, *Nucl. Phys. B* **234** (1984) 173 [[INSPIRE](#)].
- [50] C. Vafa and E. Witten, *Parity conservation in QCD*, *Phys. Rev. Lett.* **53** (1984) 535 [[INSPIRE](#)].
- [51] P. Di Vecchia and G. Veneziano, *Chiral dynamics in the large N limit*, *Nucl. Phys. B* **171** (1980) 253 [[INSPIRE](#)].
- [52] G. Grilli di Cortona, E. Hardy, J. Pardo Vega and G. Villadoro, *The QCD axion, precisely*, *JHEP* **01** (2016) 034 [[arXiv:1511.02867](#)] [[INSPIRE](#)].
- [53] S. Borsányi et al., *Axion cosmology, lattice QCD and the dilute instanton gas*, *Phys. Lett. B* **752** (2016) 175 [[arXiv:1508.06917](#)] [[INSPIRE](#)].
- [54] S. Borsányi et al., *Calculation of the axion mass based on high-temperature lattice quantum chromodynamics*, *Nature* **539** (2016) 69 [[arXiv:1606.07494](#)] [[INSPIRE](#)].
- [55] P. Petreczky, H.-P. Schadler and S. Sharma, *The topological susceptibility in finite temperature QCD and axion cosmology*, *Phys. Lett. B* **762** (2016) 498 [[arXiv:1606.03145](#)] [[INSPIRE](#)].
- [56] R.D. Pisarski and L.G. Yaffe, *The density of instantons at finite temperature*, *Phys. Lett.* **97B** (1980) 110 [[INSPIRE](#)].
- [57] J. Jaeckel and A. Ringwald, *The low-energy frontier of particle physics*, *Ann. Rev. Nucl. Part. Sci.* **60** (2010) 405 [[arXiv:1002.0329](#)] [[INSPIRE](#)].
- [58] J.L. Feng, *Dark matter candidates from particle physics and methods of detection*, *Ann. Rev. Astron. Astrophys.* **48** (2010) 495 [[arXiv:1003.0904](#)] [[INSPIRE](#)].
- [59] S. Weinberg, *Cosmology*, Oxford University Press, Oxford U.K. (2008).

- [60] J.A. Frieman and A.H. Jaffe, *Cosmological constraints on pseudoNambu-Goldstone bosons*, *Phys. Rev. D* **45** (1992) 2674 [[INSPIRE](#)].
- [61] D.H. Lyth, *Axions and inflation: sitting in the vacuum*, *Phys. Rev. D* **45** (1992) 3394 [[INSPIRE](#)].
- [62] K. Strobl and T.J. Weiler, *Anharmonic evolution of the cosmic axion density spectrum*, *Phys. Rev. D* **50** (1994) 7690 [[astro-ph/9405028](#)] [[INSPIRE](#)].
- [63] K.J. Bae, J.-H. Huh and J.E. Kim, *Update of axion CDM energy*, *JCAP* **09** (2008) 005 [[arXiv:0806.0497](#)] [[INSPIRE](#)].
- [64] O. Wantz and E.P.S. Shellard, *Axion cosmology revisited*, *Phys. Rev. D* **82** (2010) 123508 [[arXiv:0910.1066](#)] [[INSPIRE](#)].
- [65] K.J. Mack and P.J. Steinhardt, *Cosmological problems with multiple axion-like fields*, *JCAP* **05** (2011) 001 [[arXiv:0911.0418](#)] [[INSPIRE](#)].
- [66] K.J. Mack, *Axions, inflation and the anthropic principle*, *JCAP* **07** (2011) 021 [[arXiv:0911.0421](#)] [[INSPIRE](#)].
- [67] T.W.B. Kibble, *Topology of cosmic domains and strings*, *J. Phys. A* **9** (1976) 1387 [[INSPIRE](#)].
- [68] E.W. Kolb and M.S. Turner, *The early universe*, Frontiers in Physics, Westview Press, U.S.A. (1994).
- [69] R.L. Davis, *Goldstone bosons in string models of galaxy formation*, *Phys. Rev. D* **32** (1985) 3172 [[INSPIRE](#)].
- [70] R.L. Davis, *Cosmic axions from cosmic strings*, *Phys. Lett. B* **180** (1986) 225 [[INSPIRE](#)].
- [71] D. Harari and P. Sikivie, *On the evolution of global strings in the early universe*, *Phys. Lett. B* **195** (1987) 361 [[INSPIRE](#)].
- [72] C. Hagmann and P. Sikivie, *Computer simulations of the motion and decay of global strings*, *Nucl. Phys. B* **363** (1991) 247 [[INSPIRE](#)].
- [73] R.A. Battye and E.P.S. Shellard, *Global string radiation*, *Nucl. Phys. B* **423** (1994) 260 [[astro-ph/9311017](#)] [[INSPIRE](#)].
- [74] R.A. Battye and E.P.S. Shellard, *Axion string constraints*, *Phys. Rev. Lett.* **73** (1994) 2954 [*Erratum ibid.* **76** (1996) 2203] [[astro-ph/9403018](#)] [[INSPIRE](#)].
- [75] C. Hagmann, S. Chang and P. Sikivie, *Axions from string decay*, *Nucl. Phys. Proc. Suppl.* **72** (1999) 81 [[hep-ph/9807428](#)] [[INSPIRE](#)].
- [76] M. Yamaguchi, M. Kawasaki and J. Yokoyama, *Evolution of axionic strings and spectrum of axions radiated from them*, *Phys. Rev. Lett.* **82** (1999) 4578 [[hep-ph/9811311](#)] [[INSPIRE](#)].
- [77] C. Hagmann, S. Chang and P. Sikivie, *Axion radiation from strings*, *Phys. Rev. D* **63** (2001) 125018 [[hep-ph/0012361](#)] [[INSPIRE](#)].
- [78] T. Hiramatsu, M. Kawasaki, K. Saikawa and T. Sekiguchi, *Production of dark matter axions from collapse of string-wall systems*, *Phys. Rev. D* **85** (2012) 105020 [*Erratum ibid.* **86** (2012) 089902] [[arXiv:1202.5851](#)] [[INSPIRE](#)].
- [79] T. Hiramatsu, M. Kawasaki and K. Saikawa, *Evolution of string-wall networks and axionic domain wall problem*, *JCAP* **08** (2011) 030 [[arXiv:1012.4558](#)] [[INSPIRE](#)].
- [80] T. Hiramatsu et al., *Improved estimation of radiated axions from cosmological axionic strings*, *Phys. Rev. D* **83** (2011) 123531 [[arXiv:1012.5502](#)] [[INSPIRE](#)].
- [81] M. Kawasaki, K. Saikawa and T. Sekiguchi, *Axion dark matter from topological defects*, *Phys. Rev. D* **91** (2015) 065014 [[arXiv:1412.0789](#)] [[INSPIRE](#)].

- [82] L. Fleury and G.D. Moore, *Axion dark matter: strings and their cores*, *JCAP* **01** (2016) 004 [[arXiv:1509.00026](#)] [[INSPIRE](#)].
- [83] V.B. Klaer and G.D. Moore, *How to simulate global cosmic strings with large string tension*, *JCAP* **10** (2017) 043 [[arXiv:1707.05566](#)] [[INSPIRE](#)].
- [84] V.B. Klaer and G.D. Moore, *The dark-matter axion mass*, *JCAP* **11** (2017) 049 [[arXiv:1708.07521](#)] [[INSPIRE](#)].
- [85] M. Gorghetto, E. Hardy and G. Villadoro, *Axions from strings: the attractive solution*, *JHEP* **07** (2018) 151 [[arXiv:1806.04677](#)] [[INSPIRE](#)].
- [86] P. Sikivie, *Axion cosmology*, *Lect. Notes Phys.* **741** (2008) 19 [[astro-ph/0610440](#)] [[INSPIRE](#)].
- [87] M.S. Turner, *Early-universe thermal production of not-so-invisible axions*, *Phys. Rev. Lett.* **59** (1987) 2489 [Erratum *ibid.* **60** (1988) 1101] [[INSPIRE](#)].
- [88] S. Chang and K. Choi, *Hadronic axion window and the big bang nucleosynthesis*, *Phys. Lett. B* **316** (1993) 51 [[hep-ph/9306216](#)] [[INSPIRE](#)].
- [89] S. Hannestad, A. Mirizzi and G. Raffelt, *A new cosmological mass limit on thermal relic axions*, *JCAP* **07** (2005) 002 [[hep-ph/0504059](#)] [[INSPIRE](#)].
- [90] E. Masso, F. Rota and G. Zsembinszki, *Axion thermalization in the early universe*, *Phys. Rev. D* **66** (2002) 023004 [[hep-ph/0203221](#)] [[INSPIRE](#)].
- [91] P. Graf and F.D. Steffen, *Thermal axion production in the primordial quark-gluon plasma*, *Phys. Rev. D* **83** (2011) 075011 [[arXiv:1008.4528](#)] [[INSPIRE](#)].
- [92] A. Salvio, A. Strumia and W. Xue, *Thermal axion production*, *JCAP* **01** (2014) 011 [[arXiv:1310.6982](#)] [[INSPIRE](#)].
- [93] R.Z. Ferreira and A. Notari, *Observable windows for the QCD axion through the number of relativistic species*, *Phys. Rev. Lett.* **120** (2018) 191301 [[arXiv:1801.06090](#)] [[INSPIRE](#)].
- [94] S. Hannestad and G. Raffelt, *Cosmological mass limits on neutrinos, axions and other light particles*, *JCAP* **04** (2004) 008 [[hep-ph/0312154](#)] [[INSPIRE](#)].
- [95] A. Melchiorri, O. Mena and A. Slosar, *Improved cosmological bound on the thermal axion mass*, *Phys. Rev. D* **76** (2007) 041303 [[arXiv:0705.2695](#)] [[INSPIRE](#)].
- [96] S. Hannestad, A. Mirizzi, G.G. Raffelt and Y.Y.Y. Wong, *Cosmological constraints on neutrino plus axion hot dark matter*, *JCAP* **08** (2007) 015 [[arXiv:0706.4198](#)] [[INSPIRE](#)].
- [97] S. Hannestad, A. Mirizzi, G.G. Raffelt and Y.Y.Y. Wong, *Cosmological constraints on neutrino plus axion hot dark matter: Update after WMAP-5*, *JCAP* **04** (2008) 019 [[arXiv:0803.1585](#)] [[INSPIRE](#)].
- [98] S. Hannestad, A. Mirizzi, G.G. Raffelt and Y.Y.Y. Wong, *Neutrino and axion hot dark matter bounds after WMAP-7*, *JCAP* **08** (2010) 001 [[arXiv:1004.0695](#)] [[INSPIRE](#)].
- [99] M. Archidiacono et al., *Axion hot dark matter bounds after Planck*, *JCAP* **10** (2013) 020 [[arXiv:1307.0615](#)] [[INSPIRE](#)].
- [100] E. Giusarma et al., *Relic neutrinos, thermal axions and cosmology in early 2014*, *Phys. Rev. D* **90** (2014) 043507 [[arXiv:1403.4852](#)] [[INSPIRE](#)].
- [101] E. Di Valentino, S. Gariazzo, E. Giusarma and O. Mena, *Robustness of cosmological axion mass limits*, *Phys. Rev. D* **91** (2015) 123505 [[arXiv:1503.00911](#)] [[INSPIRE](#)].
- [102] E. Di Valentino et al., *Cosmological axion and neutrino mass constraints from Planck 2015 temperature and polarization data*, *Phys. Lett. B* **752** (2016) 182 [[arXiv:1507.08665](#)] [[INSPIRE](#)].

- [103] D. Grin, T.L. Smith and M. Kamionkowski, *Axion constraints in non-standard thermal histories*, *Phys. Rev. D* **77** (2008) 085020 [[arXiv:0711.1352](#)] [[INSPIRE](#)].
- [104] D.B. Kaplan, *Opening the axion window*, *Nucl. Phys. B* **260** (1985) 215 [[INSPIRE](#)].
- [105] M. Srednicki, *Axion couplings to matter. 1. CP conserving parts*, *Nucl. Phys. B* **260** (1985) 689 [[INSPIRE](#)].
- [106] H. Georgi, D.B. Kaplan and L. Randall, *Manifesting the invisible axion at low-energies*, *Phys. Lett.* **169B** (1986) 73 [[INSPIRE](#)].
- [107] M.J. Dolan, F. Kahlhoefer, C. McCabe and K. Schmidt-Hoberg, *A taste of dark matter: flavour constraints on pseudoscalar mediators*, *JHEP* **03** (2015) 171 [Erratum *ibid.* **07** (2015) 103] [[arXiv:1412.5174](#)] [[INSPIRE](#)].
- [108] F. Wilczek, *Axions and family symmetry breaking*, *Phys. Rev. Lett.* **49** (1982) 1549 [[INSPIRE](#)].
- [109] Y. Ema, K. Hamaguchi, T. Moroi and K. Nakayama, *Flaxion: a minimal extension to solve puzzles in the standard model*, *JHEP* **01** (2017) 096 [[arXiv:1612.05492](#)] [[INSPIRE](#)].
- [110] L. Calibbi et al., *Minimal axion model from flavor*, *Phys. Rev. D* **95** (2017) 095009 [[arXiv:1612.08040](#)] [[INSPIRE](#)].
- [111] M. Czerny and F. Takahashi, *Multi-natural inflation*, *Phys. Lett. B* **733** (2014) 241 [[arXiv:1401.5212](#)] [[INSPIRE](#)].
- [112] S.L. Cheng, C.Q. Geng and W.T. Ni, *Axion-photon couplings in invisible axion models*, *Phys. Rev. D* **52** (1995) 3132 [[hep-ph/9506295](#)] [[INSPIRE](#)].
- [113] J.E. Kim, *Constraints on very light axions from cavity experiments*, *Phys. Rev. D* **58** (1998) 055006 [[hep-ph/9802220](#)] [[INSPIRE](#)].
- [114] *Particle Data Group, Review of particle physics 2017 update* (2017).
- [115] L. Di Luzio, F. Mescia and E. Nardi, *Redefining the axion window*, *Phys. Rev. Lett.* **118** (2017) 031801 [[arXiv:1610.07593](#)] [[INSPIRE](#)].
- [116] L. Di Luzio, F. Mescia and E. Nardi, *Window for preferred axion models*, *Phys. Rev. D* **96** (2017) 075003 [[arXiv:1705.05370](#)] [[INSPIRE](#)].
- [117] P. Agrawal, J. Fan, M. Reece and L.-T. Wang, *Experimental targets for photon couplings of the QCD axion*, *JHEP* **02** (2018) 006 [[arXiv:1709.06085](#)] [[INSPIRE](#)].
- [118] L. Di Luzio et al., *Astrophobic axions*, *Phys. Rev. Lett.* **120** (2018) 261803 [[arXiv:1712.04940](#)] [[INSPIRE](#)].
- [119] E. Armengaud et al., *Axion searches with the EDELWEISS-II experiment*, *JCAP* **11** (2013) 067 [[arXiv:1307.1488](#)] [[INSPIRE](#)].
- [120] K. Ehret et al., *New ALPS results on hidden-sector lightweights*, *Phys. Lett. B* **689** (2010) 149 [[arXiv:1004.1313](#)] [[INSPIRE](#)].
- [121] CAST collaboration, *An improved limit on the axion-photon coupling from the CAST experiment*, *JCAP* **04** (2007) 010 [[hep-ex/0702006](#)] [[INSPIRE](#)].
- [122] CAST collaboration, *New CAST limit on the axion-photon interaction*, *Nature Phys.* **13** (2017) 584 [[arXiv:1705.02290](#)] [[INSPIRE](#)].
- [123] S. De Panfilis et al., *Limits on the abundance and coupling of cosmic axions at $4.5\,\mu\text{ev} < m_a < 5.0\,\mu\text{ev}$* , *Phys. Rev. Lett.* **59** (1987) 839 [[INSPIRE](#)].
- [124] W. Wuensch et al., *Results of a laboratory search for cosmic axions and other weakly coupled light particles*, *Phys. Rev. D* **40** (1989) 3153 [[INSPIRE](#)].

- [125] C. Hagmann, P. Sikivie, N.S. Sullivan and D.B. Tanner, *Results from a search for cosmic axions*, *Phys. Rev. D* **42** (1990) 1297 [[INSPIRE](#)].
- [126] ADMX collaboration, *Results from a high sensitivity search for cosmic axions*, *Phys. Rev. Lett.* **80** (1998) 2043 [[astro-ph/9801286](#)] [[INSPIRE](#)].
- [127] ADMX collaboration, *Large scale microwave cavity search for dark matter axions*, *Phys. Rev. D* **64** (2001) 092003 [[INSPIRE](#)].
- [128] ADMX collaboration, *An improved RF cavity search for halo axions*, *Phys. Rev. D* **69** (2004) 011101 [[astro-ph/0310042](#)] [[INSPIRE](#)].
- [129] ADMX collaboration, *A high resolution search for dark-matter axions*, *Phys. Rev. D* **74** (2006) 012006 [[astro-ph/0603108](#)] [[INSPIRE](#)].
- [130] ADMX collaboration, *A SQUID-based microwave cavity search for dark-matter axions*, *Phys. Rev. Lett.* **104** (2010) 041301 [[arXiv:0910.5914](#)] [[INSPIRE](#)].
- [131] ADMX collaboration, *A search for invisible axion dark matter with the axion dark matter experiment*, *Phys. Rev. Lett.* **120** (2018) 151301 [[arXiv:1804.05750](#)] [[INSPIRE](#)].
- [132] PLANCK collaboration, *Planck 2015 results. XIII. Cosmological parameters*, *Astron. Astrophys.* **594** (2016) A13 [[arXiv:1502.01589](#)] [[INSPIRE](#)].
- [133] H.E.S.S. collaboration, *Constraints on axionlike particles with H.E.S.S. from the irregularity of the PKS 2155-304 energy spectrum*, *Phys. Rev. D* **88** (2013) 102003 [[arXiv:1311.3148](#)] [[INSPIRE](#)].
- [134] A. Payez et al., *Revisiting the SN1987A gamma-ray limit on ultralight axion-like particles*, *JCAP* **02** (2015) 006 [[arXiv:1410.3747](#)] [[INSPIRE](#)].
- [135] A. Ayala et al., *Revisiting the bound on axion-photon coupling from globular clusters*, *Phys. Rev. Lett.* **113** (2014) 191302 [[arXiv:1406.6053](#)] [[INSPIRE](#)].
- [136] K. Van Bibber et al., *Proposed experiment to produce and detect light pseudoscalars*, *Phys. Rev. Lett.* **59** (1987) 759 [[INSPIRE](#)].
- [137] A.A. Anselm, *Arion $\leftrightarrow$ photon oscillations in a steady magnetic field* (in Russian), *Yad. Fiz.* **42** (1985) 1480 [[INSPIRE](#)].
- [138] A.A. Anselm, *Experimental test for arion $\leftrightarrow$ photon oscillations in a homogeneous constant magnetic field*, *Phys. Rev. D* **37** (1988) 2001 [[INSPIRE](#)].
- [139] ALPS collaboration, *Resonant laser power build-up in ALPS: a ‘Light-shining-through-walls’ experiment*, *Nucl. Instrum. Meth. A* **612** (2009) 83 [[arXiv:0905.4159](#)] [[INSPIRE](#)].
- [140] OSQAR collaboration, *First results from the OSQAR photon regeneration experiment: no light shining through a wall*, *Phys. Rev. D* **78** (2008) 092003 [[arXiv:0712.3362](#)] [[INSPIRE](#)].
- [141] OSQAR collaboration, *New exclusion limits on scalar and pseudoscalar axionlike particles from light shining through a wall*, *Phys. Rev. D* **92** (2015) 092002 [[arXiv:1506.08082](#)] [[INSPIRE](#)].
- [142] G.J. Feldman and R.D. Cousins, *A unified approach to the classical statistical analysis of small signals*, *Phys. Rev. D* **57** (1998) 3873 [[physics/9711021](#)] [[INSPIRE](#)].
- [143] P. Sikivie, *Experimental tests of the invisible axion*, *Phys. Rev. Lett.* **51** (1983) 1415 [*Erratum ibid.* **52** (1984) 695] [[INSPIRE](#)].
- [144] P. Sikivie, *Detection rates for “invisible”-axion searches*, *Phys. Rev. D* **32** (1985) 2988.
- [145] 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 [[INSPIRE](#)].

- [146] J. Redondo, *Solar axion flux from the axion-electron coupling*, *JCAP* **12** (2013) 008 [[arXiv:1310.0823](#)] [[INSPIRE](#)].
- [147] H. Primakoff, *Photo-Production of neutral mesons in nuclear electric fields and the mean life of the neutral meson*, *Phys. Rev.* **81** (1951) 899.
- [148] A. Halprin, C.M. Andersen and H. Primakoff, *Photonic decay rates and nuclear-Coulomb-field coherent production processes*, *Phys. Rev.* **152** (1966) 1295 [[INSPIRE](#)].
- [149] G.G. Raffelt, *Astrophysical axion bounds diminished by screening effects*, *Phys. Rev. D* **33** (1986) 897 [[INSPIRE](#)].
- [150] G.G. Raffelt, *Plasmon decay into low mass bosons in stars*, *Phys. Rev. D* **37** (1988) 1356 [[INSPIRE](#)].
- [151] S. Dimopoulos, J.A. Frieman, B.W. Lynn and G.D. Starkman, *Axiorecombination: a new mechanism for stellar axion production*, *Phys. Lett. B* **179** (1986) 223 [[INSPIRE](#)].
- [152] S. Dimopoulos, G.D. Starkman and B.W. Lynn, *Atomic enhancements in the detection of axions*, *Mod. Phys. Lett. A* **1** (1986) 491 [[INSPIRE](#)].
- [153] M. Pospelov, A. Ritz and M.B. Voloshin, *Bosonic super-WIMPs as keV-scale dark matter*, *Phys. Rev. D* **78** (2008) 115012 [[arXiv:0807.3279](#)] [[INSPIRE](#)].
- [154] A. Derevianko, V.A. Dzuba, V.V. Flambaum and M. Pospelov, *Axio-electric effect*, *Phys. Rev. D* **82** (2010) 065006 [[arXiv:1007.1833](#)] [[INSPIRE](#)].
- [155] A.R. Zhitnitsky and Yu. I. Skovpen, *On production and detecting of axions at penetration of electrons through matter* (in Russian), *Sov. J. Nucl. Phys.* **29** (1979) 513 [*Yad. Fiz.* **29** (1979) 995] [[INSPIRE](#)].
- [156] L.M. Krauss, J.E. Moody and F. Wilczek, *A stellar energy loss mechanism involving axions*, *Phys. Lett.* **144B** (1984) 391 [[INSPIRE](#)].
- [157] K.O. Mikaelian, *Astrophysical implications of new light Higgs bosons*, *Phys. Rev. D* **18** (1978) 3605 [[INSPIRE](#)].
- [158] M. Fukugita, S. Watamura and M. Yoshimura, *Light pseudoscalar particle and stellar energy loss*, *Phys. Rev. Lett.* **48** (1982) 1522 [[INSPIRE](#)].
- [159] M. Fukugita, S. Watamura and M. Yoshimura, *Astrophysical constraints on a new light axion and other weakly interacting particles*, *Phys. Rev. D* **26** (1982) 1840 [[INSPIRE](#)].
- [160] A. Serenelli, S. Basu, J.W. Ferguson and M. Asplund, *New solar composition: the problem with solar models revisited*, *Astrophys. J.* **705** (2009) L123 [[arXiv:0909.2668](#)] [[INSPIRE](#)].
- [161] M. Asplund, N. Grevesse, A.J. Sauval and P. Scott, *The chemical composition of the Sun*, *Ann. Rev. Astron. Astrophys.* **47** (2009) 481 [[arXiv:0909.0948](#)] [[INSPIRE](#)].
- [162] N. Vinyoles et al., *A new generation of standard solar models*, *Astrophys. J.* **835** (2017) 202 [[arXiv:1611.09867](#)] [[INSPIRE](#)].
- [163] N. Grevesse and A.J. Sauval, *Standard solar composition*, *Space Sci. Rev.* **85** (1998) 161.
- [164] 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](#)] [[INSPIRE](#)].
- [165] K. Barth et al., *CAST constraints on the axion-electron coupling*, *JCAP* **05** (2013) 010 [[arXiv:1302.6283](#)] [[INSPIRE](#)].
- [166] S. Turck-Chieze et al., *Solar neutrino emission deduced from a seismic model*, *Astrophys. J.* **555** (2001) L69 [[INSPIRE](#)].
- [167] J. Hoskins et al., *A search for non-virialized axionic dark matter*, *Phys. Rev. D* **84** (2011) 121302 [[arXiv:1109.4128](#)] [[INSPIRE](#)].

- [168] L. Krauss, J. Moody, F. Wilczek and D.E. Morris, *Calculations for cosmic axion detection*, *Phys. Rev. Lett.* **55** (1985) 1797 [[INSPIRE](#)].
- [169] ADMX collaboration, *Experimental constraints on the axion dark matter halo density*, *Astrophys. J.* **571** (2002) L27 [[astro-ph/0104200](#)] [[INSPIRE](#)].
- [170] ADMX collaboration, *Design and performance of the ADMX SQUID-based microwave receiver*, *Nucl. Instrum. Meth. A* **656** (2011) 39 [[arXiv:1105.4203](#)] [[INSPIRE](#)].
- [171] L. Visinelli and P. Gondolo, *Dark matter axions revisited*, *Phys. Rev. D* **80** (2009) 035024 [[arXiv:0903.4377](#)] [[INSPIRE](#)].
- [172] L. Visinelli and P. Gondolo, *Axion cold dark matter in non-standard cosmologies*, *Phys. Rev. D* **81** (2010) 063508 [[arXiv:0912.0015](#)] [[INSPIRE](#)].
- [173] P.J. Steinhardt and M.S. Turner, *Saving the invisible axion*, *Phys. Lett.* **129B** (1983) 51 [[INSPIRE](#)].
- [174] K. Yamamoto, *Saving the axions in superstring models*, *Phys. Lett.* **161B** (1985) 289 [[INSPIRE](#)].
- [175] G. Lazarides, C. Panagiotakopoulos and Q. Shafi, *Relaxing the cosmological bound on axions*, *Phys. Lett. B* **192** (1987) 323 [[INSPIRE](#)].
- [176] D.V. Nanopoulos and M. Srednicki, *Before primordial inflation*, *Phys. Lett.* **133B** (1983) 287 [[INSPIRE](#)].
- [177] S. Dimopoulos and L.J. Hall, *Inflation and invisible axions*, *Phys. Rev. Lett.* **60** (1988) 1899 [[INSPIRE](#)].
- [178] H. Davoudiasl, D. Hooper and S.D. McDermott, *Inflatable dark matter*, *Phys. Rev. Lett.* **116** (2016) 031303 [[arXiv:1507.08660](#)] [[INSPIRE](#)].
- [179] S. Hoof and J. Jaeckel, *QCD axions and axionlike particles in a two-inflation scenario*, *Phys. Rev. D* **96** (2017) 115016 [[arXiv:1709.01090](#)] [[INSPIRE](#)].
- [180] M. Kawasaki, F. Takahashi and M. Yamada, *Suppressing the QCD axion abundance by hidden monopoles*, *Phys. Lett. B* **753** (2016) 677 [[arXiv:1511.05030](#)] [[INSPIRE](#)].
- [181] M.M. Miller Bertolami, B.E. Melendez, L.G. Althaus and J. Isern, *Revisiting the axion bounds from the galactic white dwarf luminosity function*, *JCAP* **10** (2014) 069 [[arXiv:1406.7712](#)] [[INSPIRE](#)].
- [182] J. Isern et al., *Axions and the luminosity function of white dwarfs: the thin and thick discs and the halo*, *Mon. Not. Roy. Astron. Soc.* **478** (2018) 2569 [[arXiv:1805.00135](#)] [[INSPIRE](#)].
- [183] G. Raffelt and L. Stodolsky, *Mixing of the photon with low mass particles*, *Phys. Rev. D* **37** (1988) 1237 [[INSPIRE](#)].
- [184] C. Csáki, N. Kaloper and J. Terning, *Dimming supernovae without cosmic acceleration*, *Phys. Rev. Lett.* **88** (2002) 161302 [[hep-ph/0111311](#)] [[INSPIRE](#)].
- [185] Y. Grossman, S. Roy and J. Zupan, *Effects of initial axion production and photon axion oscillation on type-IA supernova dimming*, *Phys. Lett. B* **543** (2002) 23 [[hep-ph/0204216](#)] [[INSPIRE](#)].
- [186] 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 [[arXiv:0704.3044](#)] [[INSPIRE](#)].
- [187] D. Hooper and P.D. Serpico, *Detecting axion-like particles with Gamma Ray Telescopes*, *Phys. Rev. Lett.* **99** (2007) 231102 [[arXiv:0706.3203](#)] [[INSPIRE](#)].
- [188] A. De Angelis, O. Mansutti and M. Roncadelli, *Axion-like particles, cosmic magnetic fields and gamma-ray astrophysics*, *Phys. Lett. B* **659** (2008) 847 [[arXiv:0707.2695](#)] [[INSPIRE](#)].

- [189] K.A. Hochmuth and G. Sigl, *Effects of axion-photon mixing on gamma-ray spectra from magnetized astrophysical sources*, *Phys. Rev. D* **76** (2007) 123011 [[arXiv:0708.1144](#)] [[INSPIRE](#)].
- [190] A. Mirizzi and D. Montanino, *Stochastic conversions of TeV photons into axion-like particles in extragalactic magnetic fields*, *JCAP* **12** (2009) 004 [[arXiv:0911.0015](#)] [[INSPIRE](#)].
- [191] D. Wouters and P. Brun, *Irregularity in gamma ray source spectra as a signature of axionlike particles*, *Phys. Rev. D* **86** (2012) 043005 [[arXiv:1205.6428](#)] [[INSPIRE](#)].
- [192] G. Galanti and M. Roncadelli, *Comment on “Irregularity in gamma ray source spectra as a signature of axion-like particles”*, [arXiv:1305.2114](#) [[INSPIRE](#)].
- [193] FERMI-LAT collaboration, *Search for spectral irregularities due to photon-axionlike-particle oscillations with the Fermi Large Area Telescope*, *Phys. Rev. Lett.* **116** (2016) 161101 [[arXiv:1603.06978](#)] [[INSPIRE](#)].
- [194] Z.-Q. Xia et al., *Searching for spectral oscillations due to photon-axionlike particle conversion using the Fermi-LAT observations of bright supernova remnants*, *Phys. Rev. D* **97** (2018) 063003 [[arXiv:1801.01646](#)] [[INSPIRE](#)].
- [195] J. Majumdar, F. Calore and D. Horns, *Search for gamma-ray spectral modulations in galactic pulsars*, *JCAP* **04** (2018) 048 [[arXiv:1801.08813](#)] [[INSPIRE](#)].
- [196] C. Zhang et al., *New bounds on axionlike particles from the Fermi Large Area Telescope observation of PKS 2155-304*, *Phys. Rev. D* **97** (2018) 063009 [[arXiv:1802.08420](#)] [[INSPIRE](#)].
- [197] A. Mirizzi, G.G. Raffelt and P.D. Serpico, *Photon-axion conversion in intergalactic magnetic fields and cosmological consequences*, *Lect. Notes Phys.* **741** (2008) 115 [[astro-ph/0607415](#)] [[INSPIRE](#)].
- [198] J.R. Ellis and K.A. Olive, *Constraints on light particles from supernova SN1987a*, *Phys. Lett. B* **193** (1987) 525 [[INSPIRE](#)].
- [199] G. Raffelt and D. Seckel, *Bounds on exotic particle interactions from SN 1987a*, *Phys. Rev. Lett.* **60** (1988) 1793 [[INSPIRE](#)].
- [200] M.S. Turner, *Axions from SN 1987a*, *Phys. Rev. Lett.* **60** (1988) 1797 [[INSPIRE](#)].
- [201] R. Mayle et al., *Constraints on Axions from SN 1987a*, *Phys. Lett. B* **203** (1988) 188 [[INSPIRE](#)].
- [202] J.W. Brockway, E.D. Carlson and G.G. Raffelt, *SN1987A gamma-ray limits on the conversion of pseudoscalars*, *Phys. Lett. B* **383** (1996) 439 [[astro-ph/9605197](#)] [[INSPIRE](#)].
- [203] J.A. Grifols, E. Masso and R. Toldra, *Gamma-rays from SN1987A due to pseudoscalar conversion*, *Phys. Rev. Lett.* **77** (1996) 2372 [[astro-ph/9606028](#)] [[INSPIRE](#)].
- [204] E.L. Chupp, W.T. Vestrand and C. Reppin, *Experimental limits on the radiative decay of SN1987A neutrinos*, *Phys. Rev. Lett.* **62** (1989) 505 [[INSPIRE](#)].
- [205] M.S. Pshirkov, P.G. Tinyakov, P.P. Kronberg and K.J. Newton-McGee, *Deriving global structure of the galactic magnetic field from Faraday rotation measures of extragalactic sources*, *Astrophys. J.* **738** (2011) 192 [[arXiv:1103.0814](#)] [[INSPIRE](#)].
- [206] R. Jansson and G.R. Farrar, *A new model of the galactic magnetic field*, *Astrophys. J.* **757** (2012) 14 [[arXiv:1204.3662](#)] [[INSPIRE](#)].
- [207] K. Sato and H. Sato, *Higgs meson emission from a star and a constraint on its mass*, *Prog. Theor. Phys.* **54** (1975) 1564 [[INSPIRE](#)].

- [208] G.G. Raffelt, *Astrophysical methods to constrain axions and other novel particle phenomena*, *Phys. Rept.* **198** (1990) 1.
- [209] G. G. Raffelt, *Stars as laboratories for fundamental physics: the astrophysics of neutrinos, axions, and other weakly interacting particles*, University Of Chicago Press, Chicago U.S.A. (1996).
- [210] M. Salaris, M. Riello, S. Cassisi and G. Piotto, *The initial helium abundance of the Galactic globular cluster system*, *Astron. Astrophys.* **420** (2004) 911 [[astro-ph/0403600](#)] [[INSPIRE](#)].
- [211] A. Buzzoni et al., *Helium abundance in globular clusters — The R-method*, *Astron. Astrophys.* **128** (1983) 94.
- [212] G.G. Raffelt, *Core mass at the helium flash from observations and a new bound on neutrino electromagnetic properties*, *Astrophys. J.* **365** (1990) 559 [[INSPIRE](#)].
- [213] N. Viaux et al., *Neutrino and axion bounds from the globular cluster M5 (NGC 5904)*, *Phys. Rev. Lett.* **111** (2013) 231301 [[arXiv:1311.1669](#)] [[INSPIRE](#)].
- [214] Y.I. Izotov, T.X. Thuan and N.G. Guseva, *A new determination of the primordial He abundance using the He I $\lambda 10830$ Å emission line: cosmological implications*, *Mon. Not. Roy. Astron. Soc.* **445** (2014) 778 [[arXiv:1408.6953](#)] [[INSPIRE](#)].
- [215] E. Aver, K.A. Olive and E.D. Skillman, *The effects of He I $\lambda 10830$ on helium abundance determinations*, *JCAP* **07** (2015) 011 [[arXiv:1503.08146](#)] [[INSPIRE](#)].
- [216] G.G. Raffelt, *Axion constraints from white dwarf cooling times*, *Phys. Lett.* **166B** (1986) 402 [[INSPIRE](#)].
- [217] M. Nakagawa, Y. Kohyama and N. Itoh, *Axion Bremsstrahlung in dense stars*, *Astrophys. J.* **322** (1987) 291 [[INSPIRE](#)].
- [218] M. Nakagawa, T. Adachi, Y. Kohyama and N. Itoh, *Axion bremsstrahlung in dense stars. II — Phonon contributions*, *Astrophys. J.* **326** (1988) 241 [[INSPIRE](#)].
- [219] T. Altherr, E. Petitgirard and T. del Rio Gaztelurrutia, *Axion emission from red giants and white dwarfs*, *Astropart. Phys.* **2** (1994) 175 [[hep-ph/9310304](#)] [[INSPIRE](#)].
- [220] A.H. Corsico et al., *The potential of the variable DA white dwarf G117-B15A as a tool for fundamental physics*, *New Astron.* **6** (2001) 197 [[astro-ph/0104103](#)] [[INSPIRE](#)].
- [221] 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 [[arXiv:0711.2041](#)] [[INSPIRE](#)].
- [222] J. Isern, E. Garcia-Berro, S. Torres and S. Catalan, *Axions and the cooling of white dwarf stars*, *Astrophys. J.* **682** (2008) L109 [[arXiv:0806.2807](#)] [[INSPIRE](#)].
- [223] F. Feroz and M.P. Hobson, *Multimodal nested sampling: an efficient and robust alternative to MCMC methods for astronomical data analysis*, *Mon. Not. Roy. Astron. Soc.* **384** (2008) 449 [[arXiv:0704.3704](#)] [[INSPIRE](#)].
- [224] F. Feroz, M.P. Hobson and M. Bridges, *MultiNest: an efficient and robust Bayesian inference tool for cosmology and particle physics*, *Mon. Not. Roy. Astron. Soc.* **398** (2009) 1601 [[arXiv:0809.3437](#)] [[INSPIRE](#)].
- [225] F. Feroz, M.P. Hobson, E. Cameron and A.N. Pettitt, *Importance nested sampling and the MultiNest algorithm*, [arXiv:1306.2144](#) [[INSPIRE](#)].
- [226] P. Svrček and E. Witten, *Axions in string theory*, *JHEP* **06** (2006) 051 [[hep-th/0605206](#)] [[INSPIRE](#)].
- [227] L. Visinelli and S. Vagnozzi, *Cosmological window onto the string axiverse and the supersymmetry breaking scale*, [arXiv:1809.06382](#) [[INSPIRE](#)].

- [228] C.-Y. Chen and S. Dawson, *Exploring two Higgs doublet models through Higgs production*, *Phys. Rev. D* **87** (2013) 055016 [[arXiv:1301.0309](#)] [[INSPIRE](#)].
- [229] H. Jeffreys, *The theory of probability*, Oxford Classic Texts in the Physical Sciences. Oxford University Press, Oxford U.K. (1939).
- [230] R.E. Kass and A.E. Raftery, *Bayes factors*, *J. Amer. Stat. Ass.* **90** (1995) 773.
- [231] J.O. Berger et al., *Objective Bayesian methods for model selection: introduction and comparison*, *Lecture Notes-Monograph Series* **38** (2001) 135.
- [232] J. Meija et al., *Atomic weights of the elements 2013 (IUPAC Technical Report)*, *Pure Appl. Chem.* **88** (2015) 265.
- [233] PARTICLE DATA GROUP collaboration, *Review of particle physics*, *Chin. Phys. C* **40** (2016) 100001 [[INSPIRE](#)].