

Cosmological constraints on decaying axion-like particles: a global analysis

Csaba Balázs,¹ Sanjay Bloor,² Tomás E. Gonzalo,^{3,4} Will Handley,^{5,6} Sebastian Hoof,⁷ Felix Kahlhoefer,^{3,4} Marie Lecroq,^{1,8} David J. E. Marsh,⁹ Janina Renk,^{2,10,11} Pat Scott^{2,10} and Patrick Stöcker^{3,12}

¹School of Physics and Astronomy, Monash University, Melbourne, VIC 3800, Australia

²Department of Physics, Imperial College London, Blackett Laboratory, Prince Consort Road, London SW7 2AZ, UK

³Institute for Theoretical Particle Physics and Cosmology (TTK), RWTH Aachen University, D-52056 Aachen, Germany

⁴Institute for Theoretical Particle Physics (TTP), Karlsruhe Institute of Technology (KIT), 76128 Karlsruhe, Germany

⁵Cavendish Laboratory, University of Cambridge, JJ Thomson Avenue, Cambridge, CB3 0HE, UK

⁶Kavli Institute for Cosmology, University of Cambridge, Madingley Road, Cambridge, CB3 0HA, UK

⁷Institut für Astrophysik und Geophysik, Georg-August-Universität Göttingen, Friedrich-Hund-Platz 1, 37077 Göttingen, Germany

⁸Institut d'Astrophysique de Paris, 98 bis boulevard Arago, 75014 Paris, France

⁹Theoretical Particle Physics and Cosmology Group, Department of Physics, King's College London, Strand, London, WC2R 2LS, UK

¹⁰Oskar Klein Centre for Cosmoparticle Physics, AlbaNova University Centre, SE-10691 Stockholm, Sweden

¹¹School of Mathematics and Physics, The University of Queensland, St. Lucia, Brisbane, QLD 4072, Australia

¹²Institute for Astroparticle Physics (IAP), Karlsruhe Institute of Technology (KIT), Hermann-von-Helmholtz-Platz 1, D-76344 Eggenstein-Leopoldshafen, Germany

E-mail: tomas.gonzalo@kit.edu, hoof@uni-goettingen.de, kahlhoefer@kit.edu,
stoecker@physik.rwth-aachen.de

Abstract. Axion-like particles (ALPs) decaying into photons are known to affect a wide range of astrophysical and cosmological observables. In this study we focus on ALPs with masses in the keV–MeV range and lifetimes between 10^4 and 10^{13} seconds, corresponding to decays between the end of Big Bang Nucleosynthesis and the formation of the Cosmic Microwave Background (CMB). Using the `CosmoBit` module of the global fitting framework `GAMBIT`, we combine state-of-the-art calculations of the irreducible ALP freeze-in abundance, primordial element abundances (including photodisintegration through ALP decays), CMB spectral distortions and anisotropies, and constraints from supernovae and stellar cooling. This approach makes it possible for the first time to perform a global analysis of the ALP parameter space while varying the parameters of Λ CDM as well as several nuisance parameters. We find a lower bound on the ALP mass of around $m_a > 300$ keV, which can only be evaded if ALPs are stable on cosmological timescales. Future observations of CMB spectral distortions with a PIXIE-like mission are expected to improve this bound by two orders of magnitude.

Contents

1	Introduction	1
2	Production and decay of axion-like particles coupled to photons	4
3	Constraints and likelihoods	7
3.1	CMB power spectra	7
3.1.1	Modifications of N_{eff} and η_b	8
3.2	Primordial nucleosynthesis	10
3.2.1	Likelihood implementation and uncertainties	11
3.2.2	The ${}^7\text{Li}$ problem	12
3.3	Spectral distortions	12
3.3.1	ALP distortion signal	13
3.3.2	Likelihood implementation	15
3.4	ALP decays after SN1987A	15
3.5	Stellar evolution and cluster counts	16
3.6	Other likelihoods	17
4	Results	17
4.1	Frequentist scans	17
4.2	Bayesian scans	22
5	Conclusions	25
A	GAMBIT implementation	27
A.1	New models	27
A.2	New capabilities provided in CosmoBit and DarkBit	29
A.3	New and updated backend interfaces	31
B	Details of implementation of entropy injection between BBN and recombination	32
B.1	Numerical implementation	32
C	Details of the SN1987A limits from ALP decay	33
C.1	Reference frame transformations	33
C.2	Monte Carlo procedure	34
C.3	Results	35
C.4	Technical details	35

1 Introduction

Cosmological observables such as the Cosmic Microwave Background (CMB) have now been measured with such precision that they not only constrain the properties of particles in the primordial thermal plasma, but also even probe out-of-equilibrium species with sub-dominant cosmological abundances. For example, unstable particles that decay around the time of recombination leave an observable imprint on the CMB temperature anisotropies, even if

they change the energy density of the electron-photon plasma only at the level of one part per billion [1–3]. Earlier decays are in turn tightly constrained by observations of CMB spectral distortions [4, 5] and the primordial element abundances produced in Big Bang Nucleosynthesis (BBN) [6–8]. This raises the exciting possibility of using cosmology to search for the effects of decaying exotic particles across many different cosmological eras, ranging from one second to billions of years [9].

A particularly well-motivated class of exotic particles that may be probed in this way are axion-like particles (ALPs) [10, 11], which generically arise in theories with spontaneously broken global chiral symmetries and theories with compactified extra dimensions [12]. In analogy with the QCD axion proposed to solve the strong CP problem [13, 14], ALPs would have an approximate shift symmetry that explains both the smallness of the ALP mass and the weakness of its couplings in terms of the large scale of symmetry breaking. Although ALPs may in general have a wide range of different couplings [15, 16], the most phenomenologically interesting one is the coupling to photons. Indeed, there is a vast literature on how to constrain ALP-photon couplings using cosmological [17–20], astrophysical [21–23] and laboratory data [24, 25]. For more details, we refer the interested reader to reviews on QCD axion and ALP cosmology [26] and experimental search strategies [27], as well as a continuously updated online repository for limits and sensitivity forecasts [28].

Most studies of cosmological ALP constraints make one of two assumptions on the ALP abundance. Either ALPs are assumed to be produced by the realignment mechanism in such a way that their present-day abundance is consistent with the observed abundance of dark matter (DM) in the universe, or they are assumed to be in thermal equilibrium with the bath of Standard Model (SM) particles at some high temperature but then decay, depleting their abundance. Both of these scenarios turn out to be very tightly constrained, such that ALPs are often believed to be excluded unless they are either very light (with a mass at or below the keV-scale) or very heavy (with a mass at or above the GeV-scale) [19]. It has only recently been pointed out that a low reheating temperature can open up otherwise ruled-out parameter space for intermediate-mass ALPs, if realignment production is inefficient and their couplings are small enough to avoid attaining thermal equilibrium [29]. In this case the dominant contribution to ALP production arises from the freeze-in mechanism at temperatures close to the reheating temperature, and the resulting abundance may be small enough to evade cosmological constraints arising from their subsequent decays [19, 30].

In the present work we perform a detailed analysis of ALPs in the keV–GeV mass range with tiny couplings to SM photons, such that they never enter into thermal equilibrium in the early universe. While these ALPs may in principle be sufficiently long-lived to survive until the late universe, we focus on ALPs that decay between BBN and recombination. In this case the ALP decays affect a wide range of observables, such as

- the primordial element abundances through photodisintegration of the nuclei formed during BBN,
- the black-body spectrum of the CMB through the injection of electromagnetic energy into the plasma,
- the CMB anisotropies through modifications of the recombination history,
- the ratio of the photon and neutrino temperatures (encoded in the relativistic number of effective degrees of freedom) through a relative increase of the former,

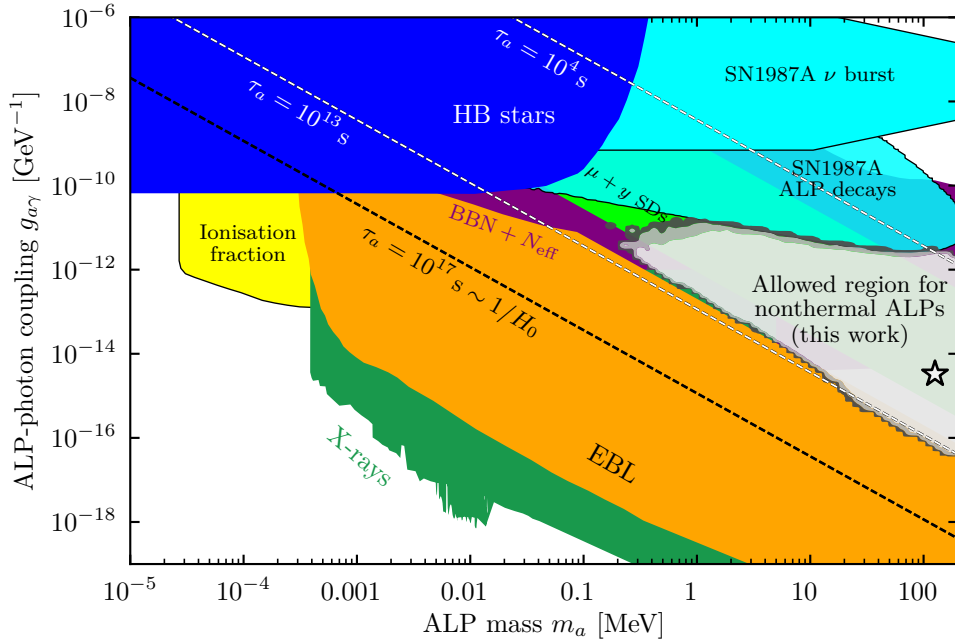


Figure 1: Summary plot of our results in the m_a - $g_{a\gamma}$ plane. We show a part of the cosmological ALP region commonly excluded by Horizontal Branch (HB) stars from cluster counts [34], supernova SN1987A constraints (see Ref. [17] and Section 3.4), ionisation fraction constraints [19], BBN and N_{eff} [29], spectral distortions [19], comparison to the extragalactic background light (EBL) [19], and X-ray searches [19, 35–37] (we used data provided by Ref. [28] for some of the contours). For nonthermal ALPs considered in this work, the grey-shaded region around the best-fitting point (star) is however not excluded.

- auxiliary astrophysical observables, such as axion emission and decays from supernova SN1987A or stellar cooling.

We make use of the recently released CosmoBit module [31] of the GAMBIT framework [32, 33] to perform a combined analysis of all these observables. Exploring the combined likelihood function has the advantage of automatically incorporating any complementary information contained in the data. Similar analyses, such as the ones presented in Refs. [19, 20], investigated the individual likelihoods instead. Compared to Ref. [19] we include additional constraints from SN1987A and a more rigorous treatment of spectral distortions and CMB constraints. In contrast to our work, Ref. [20] focuses on models with lifetimes $\tau_a \lesssim 10^7$ s and μ -type distortions only. In addition, our analysis improves upon the results of Ref. [29] by performing a more accurate calculation of the ALP abundance as a function of the reheating temperature and by including state-of-the-art likelihoods for all cosmological data sets – in particular CMB spectral distortions and anisotropies.

A visual summary of our findings in the mass-coupling plane is shown in Fig. 1. We identify large regions of viable parameter space where ALPs would constitute a sizeable fraction of DM in the very early universe, but decay away before recombination. Leaving the reheating temperature as a free parameter (imposing only the requirement $T_{\text{R}} \gtrsim 5$ MeV to ensure that BBN proceeds as usual) opens up parameter regions that are usually considered to be excluded by cosmology. Moreover, by combining the individual observations into a global likelihood function we find that ALP decays can in fact improve the agreement between

BBN and the CMB compared to Λ CDM. Intriguingly, the corresponding parameter regions may be fully probed by future missions that measure CMB spectral distortions.

The remainder of this work is structured as follows. In Section 2 we discuss the freeze-in production and subsequent decay of ALPs with tiny couplings to photons. The various ways in which these couplings can be constrained through observations and the corresponding likelihood functions are presented in Section 3. Section 4 contains the setup of our parameter scans and the results that we obtain. These results, along with a GAMBIT input file to reproduce them, and a plotting script, can be downloaded from Zenodo [38].

2 Production and decay of axion-like particles coupled to photons

In this work we will focus on ALPs a with an effective coupling to photons given by

$$\mathcal{L} = \frac{g_{a\gamma}}{4} a F^{\mu\nu} \tilde{F}_{\mu\nu}, \quad (2.1)$$

where $F^{\mu\nu}$ is the electromagnetic (EM) field strength tensor, $\tilde{F}^{\mu\nu} = \frac{1}{2}\epsilon^{\mu\nu\rho\sigma}F_{\rho\sigma}$ its dual, and the ALP-photon coupling $g_{a\gamma}$ has mass dimension -1 .

The ALP lifetime. Axion-like particles are unstable and decay into two photons with a lifetime of $\tau_a = \Gamma_a^{-1}$, where

$$\Gamma_a = \frac{g_{a\gamma}^2 m_a^3}{64\pi} = \frac{1}{1.32 \times 10^8 \text{ s}} \left(\frac{g_{a\gamma}}{10^{-12} \text{ GeV}^{-1}} \right)^2 \left(\frac{m_a}{10 \text{ MeV}} \right)^3. \quad (2.2)$$

We can broadly distinguish three different lifetime regimes:

1. ALPs decay before neutrino decoupling ($\tau_a \lesssim 1 \text{ s}$). In this case the decay photons will fully thermalise and not affect any cosmological observables.
2. ALPs decay between neutrino decoupling and the present day ($1 \text{ s} \lesssim \tau_a \lesssim 10^{17} \text{ s}$). In this case ALPs may significantly affect various cosmological observables.
3. ALPs survive until the present day ($\tau_a \gtrsim 10^{17} \text{ s}$). In this case ALPs contribute to the DM abundance; constraints in the mass range of interest arise from searches for mono-energetic X-ray or γ -ray photons.

The second case is particularly interesting because non-trivial constraints on the ALP parameter space can be expected. Indeed, the previous studies mentioned in Section 1 showed that there are very tight constraints on EM decays of DM particles during or after recombination, or during BBN. Less attention has been given to the case of ALPs decaying between BBN and recombination, i.e. $10^4 \text{ s} < \tau_a < 10^{13} \text{ s}$, which will be the focus of this work.

The ALP abundance. In addition to m_a and $g_{a\gamma}$, the two model parameters introduced above, the ALP abundance can be introduced as a third parameter. The abundance will be decisive for the subsequent analysis. Since in our set-up there are no processes that can produce ALPs in the temperature range that we will be interested in,¹ the comoving number density $a^3 n_a$ only changes through ALP decays. We can therefore define the would-be

¹A possible exception would be the inverse decay $\gamma\gamma \rightarrow a$. This process however turns out to be irrelevant in the viable parameter space of the model.

present-day abundance in the absence of decays $n_{a,0}^{\tau_a \rightarrow \infty}$ and the corresponding would-be dark matter fraction

$$\xi = \frac{m_a n_{a,0}^{\tau_a \rightarrow \infty}}{\rho_{\text{DM},0}}, \quad (2.3)$$

where $\rho_{\text{DM},0}$ is the present-day DM density. Because the comoving DM density is constant, the ALP fraction is temperature-independent in the absence of decays. When including ALP decays, however, the actual present-day abundance of ALPs becomes completely negligible for the lifetimes considered in this work. Hence, ξ is not directly constrained by the observed DM abundance, and it is possible to have $\xi > 1$.

To study the cosmological effects of ALPs at the keV–MeV scale, Ref. [19] assumes that ALPs enter thermal equilibrium in the early universe and subsequently undergo thermal freeze-out. In this case, which has also been considered in Ref. [20], ξ can be calculated in terms of the other model parameters and, typically, $\xi \gg 1$.

Here we instead consider the case where ALPs’ interactions are sufficiently weak that they never thermalise. This assumption implies an upper bound on the reheating temperature T_{R} : because ALPs couple to photons through an effective interaction, the production rate of ALPs for $T \gg m_a$ scales as $\Gamma_{\text{prod}} \sim \alpha T^3 g_{a\gamma}^2$, where $\alpha \approx 1/137$ is the EM fine-structure constant. Thermalisation occurs if $\Gamma_{\text{prod}} > H$, where $H = \dot{R}/R$ is the Hubble rate in terms of the scale factor of the universe R .² During radiation domination the Hubble rate depends only on the total energy density in radiation ρ_{rad} :

$$H^2(T) \simeq \frac{\rho_{\text{rad}}}{3\bar{m}_{\text{P}}^2} = \frac{\pi^2}{90\bar{m}_{\text{P}}^2} g_*(T) T^4 \approx 0.11 g_*(T) \frac{T^4}{\bar{m}_{\text{P}}^2}, \quad (2.4)$$

where g_* denotes the number of effective relativistic degrees of freedom and $\bar{m}_{\text{P}} \equiv M_{\text{P}}/\sqrt{8\pi} \approx 2.435 \times 10^{18}$ GeV is the reduced Planck mass. Hence, to first approximation, no thermalisation occurs if $g_{a\gamma}$ satisfies

$$g_{a\gamma} \lesssim 10^{-7} \text{ GeV}^{-1} \left(\frac{5 \text{ MeV}}{T_{\text{R}}} \right)^{\frac{1}{2}}. \quad (2.5)$$

If ALPs never enter into thermal equilibrium, their abundance depends on the details of how inflation and reheating proceed, and on the entire subsequent cosmological evolution. In particular, there may be a relevant contribution to the ALP abundance from the vacuum realignment mechanism, which depends on the initial misalignment angle of the ALP field, or from the decay of cosmic strings. For all practical purposes one can therefore treat the ALP fraction ξ as a free parameter.

When doing so, it is however essential to realise that ξ cannot be made arbitrarily small. The reason is that there is an irreducible contribution to the ALP abundance from the freeze-in mechanism, i.e. the direct production of nonthermal ALPs from the thermal bath of SM particles via the Primakoff effect [19, 30]: $\gamma \ell^\pm \rightarrow a \ell^\pm$, where ℓ^\pm denotes the charged leptons.³ Let us therefore briefly review the freeze-in production of ALPs and calculate the corresponding contribution to ξ .

The freeze-in production is dominated by temperatures close to T_{R} , so that the final abundance of ALPs will be proportional to T_{R} [30, 39]. The resulting bound on T_{R} is substantially stronger than the one obtained from the requirement of nonthermalisation via

²We denote the scale factor by R instead of a to avoid confusion with the axion field.

³Analogous Primakoff processes involving charged hadrons give a negligible contribution for the temperature range that we consider.

Eq. (2.5). Conversely, as T_{R} cannot be arbitrarily low, we obtain an irreducible contribution to the abundance of ALPs in the early universe.

Let us focus on temperatures below the muon mass m_{μ} for simplicity and consider the case that the ALP mass is negligible: $m_a \ll T \ll m_{\mu} \sim 100 \text{ MeV}$. The Boltzmann equation governing the evolution of the ALP number density n_a is then given by [30]

$$\frac{dn_a}{dt} + 3Hn_a = \frac{\alpha \zeta(3) g_{a\gamma}^2 T^6}{12\pi^2} \left[2 \ln \left(\frac{T}{m_{\gamma}} \right) + 0.8194 \right], \quad (2.6)$$

where $\zeta(3) \approx 1.202$ is Apéry's constant and $m_{\gamma}^2 = 4\pi\alpha T^2/9$ is the plasmon mass for a fully ionised universe.⁴ Defining $x = T_{\text{R}}/T$ and $Y_a = n_a/s$ with $s = \frac{2\pi^2}{45} g_*^s(T) T^3$ and $g_*^s = g_* = 10.75$ in the MeV range, this becomes

$$\frac{dY_a}{dx} = 7.82 \times 10^{-5} \frac{g_{a\gamma}^2 \bar{m}_{\text{P}} T_{\text{R}}}{x^2}. \quad (2.7)$$

This equation can be readily integrated from $x = 1$ to $x \rightarrow \infty$, to yield

$$\xi = 7.82 \times 10^{-5} \frac{m_a g_{a\gamma}^2 \bar{m}_{\text{P}} T_{\text{R}} s_0}{\rho_{\text{DM},0}} = 0.022 \left(\frac{m_a}{1 \text{ MeV}} \right) \left(\frac{T_{\text{R}}}{5 \text{ MeV}} \right) \left(\frac{g_{a\gamma}}{10^{-10} \text{ GeV}^{-1}} \right)^2, \quad (2.8)$$

where $s_0 = 2891.2 \text{ cm}^{-3}$ denotes the present-day entropy density. In particular, we find that $\xi \propto T_{\text{R}}$. The requirement that BBN proceeds as usual implies the lower bound $T_{\text{R}} \gtrsim 5 \text{ MeV}$ [40], which in turn leads to a lower bound on ξ as a function of m_a and $g_{a\gamma}$.

In the derivation of Eq. (2.8) we made a few simplifications that are not necessarily valid. In particular, we treated both ALPs and electrons as massless. However, when $m_a > T_{\text{R}}$, ALP production via the Primakoff process becomes exponentially suppressed, and the lower bound on ξ is significantly relaxed. This effect can be approximately captured by multiplying Eq. (2.8) by a Boltzmann suppression factor $e^{-m_a/T_{\text{R}}}$.

To more accurately include the effect of finite masses, we also implemented the ALP model in micrOMEGAs v5.2.11. Figure 2 provides a comparison between the ALP fraction computed by multiplying Eq. (2.8) by $e^{-m_a/T_{\text{R}}}$, and that computed by micrOMEGAs. Both approaches include quantum statistics and the plasma mass of the photon. Eq. (2.8) includes a more careful treatment of the contribution of soft momenta from Ref. [30], whereas micrOMEGAs includes the sub-leading contribution from $e^+e^- \rightarrow a\gamma$. Despite these differences, Eq. (2.8) and the results from micrOMEGAs agree very well for $m_e, m_a \ll T_{\text{R}}$. At higher masses, the curves are qualitatively similar, but diverge. In order to accurately capture the suppression of ALP production for large masses, we will always use the result from micrOMEGAs in what follows. Since these calculations are computationally expensive, we perform them in advance and interpolate the tabulated results.

To conclude this section, we note that the relative contribution of ALPs to the total energy density of the universe, $\rho_a(T)/\rho_{\text{tot}}(T)$, will be largest shortly before they start decaying, i.e. for $T \approx T_{\text{dec}}$ with $H(T_{\text{dec}}) = 1/\tau_a$. Assuming that these decays happen after neutrino decoupling but before matter-radiation equality, we can estimate

$$\frac{\rho_a(T_{\text{dec}})}{\rho_{\text{tot}}(T_{\text{dec}})} = \xi \frac{\rho_{\text{DM}}(T_{\text{dec}})}{\rho_{\text{rad}}(T_{\text{dec}})} = \xi \frac{\Omega_{\text{DM}} T_{\text{dec}}}{\Omega_{\text{rad}} T_0} \sim 10^{-4} \xi \left(\frac{\tau_a}{10^4 \text{ s}} \right)^{1/2}. \quad (2.9)$$

⁴We emphasise that ALPs are typically produced with relativistic energy, such that it is not directly possible to infer the evolution of ρ_a from eq. (2.6). Only at sufficiently late times, when $T \ll m_a$, can we make use of the relation $\rho_a = m_a n_a$.

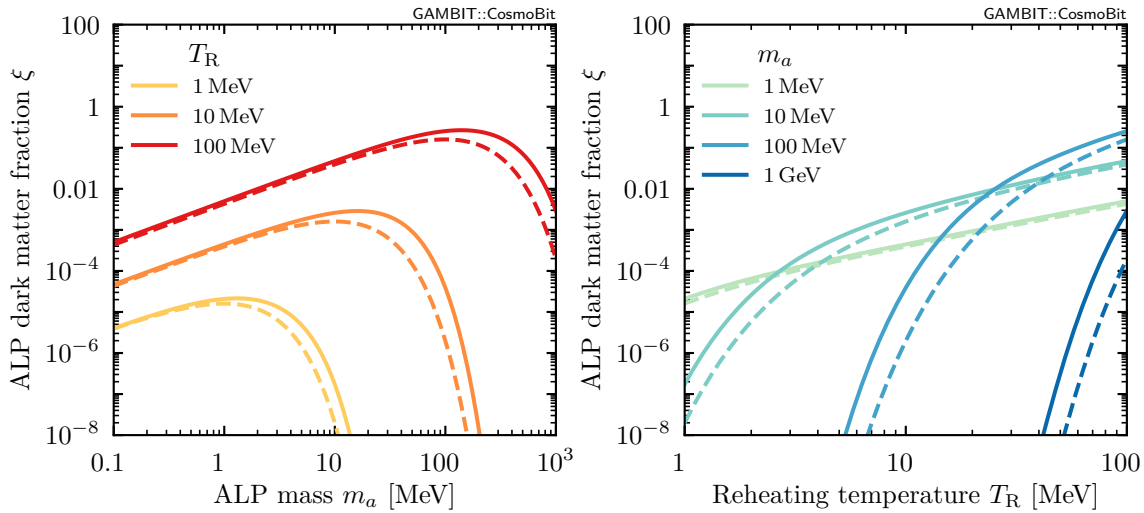


Figure 2: The irreducible freeze-in contribution to the would-be ALP dark matter fraction ξ in the absence of decays as a function of m_a (*left*) and T_R (*right*) for a reference value of $g_{a\gamma} = 10^{-11} \text{ GeV}^{-1}$. We compare the numerical result from `micrOMEGAs` v5.2.11 (solid lines) with the approximate result from Eq. (2.8) multiplied by a naïve Boltzmann suppression (dashed lines).

We will see below that in the allowed regions of parameter space this ratio is always several orders of magnitude below unity, such that ALPs never dominate the expansion history of the universe. Nevertheless, their decays can have a number of profound effects, which we will discuss in the following section.

3 Constraints and likelihoods

In this section we discuss the various constraints on the ALP model and the corresponding likelihood functions, which we implement within the `GAMBIT` software framework.

3.1 CMB power spectra

We vary all parameters of the cosmological concordance model ΛCDM : the present-day Hubble expansion rate H_0 ; the present-day baryon abundance $\omega_{b,0} = \rho_{b,0} h^2 / \rho_{\text{crit},0}$, where $H_0 = 100h \text{ km s}^{-1} \text{ Mpc}^{-1}$ and $\rho_{\text{crit},0} = 3\bar{m}_p^2 H_0^2$; the present-day dark matter abundance $\omega_{\text{DM},0}$; the redshift of reionisation z_{reio} ⁵; the amplitude A_s ; and the spectral index n_s of the initial power spectrum. We fix the temperature of the CMB⁶ to $T_{\text{CMB}} = 2.72548 \text{ K}$, the sum of the neutrino masses to $\sum m_\nu = 0.06 \text{ eV}$ and the pivot scale $k_0 = 0.05 \text{ Mpc}^{-1}$.

For a given set of parameters we use `CLASS` [41] to calculate the CMB power spectra. We then compare these predictions to data using the Planck 2018 TT,TE,EE+lowE+lensing likelihoods (referred to as the “baseline” likelihood in Ref. [42], see also Ref. [31]), which

⁵Unlike the parametrisation used in the initial `CosmoBit` release [31], τ_{reio} is replaced with z_{reio} , as the latter is the quantity that is internally used within the calculations of `CLASS` and τ_{reio} is a derived quantity. The parametrisation used here avoids numerical instabilities which can occur when τ_{reio} is used as primary parameter in scenarios of strong exotic energy injection.

⁶We note that the tiny uncertainty of the CMB temperature, $\sigma(T_{\text{CMB}}) = 5 \times 10^{-4} \text{ K}$, does become relevant in the context of CMB spectral distortions. This uncertainty is automatically marginalised over when we calculate the spectral distortion likelihood (see below).

are publicly available through `plc` [43]. For high multipoles we use the ‘lite’ version of the likelihood, which requires only one additional nuisance parameter, A_{Planck} . As done by the Planck collaboration, we apply a Gaussian prior on this parameter [43]. This is done through an additional Gaussian likelihood on this parameter, in order to allow for both, Frequentist and Bayesian analyses. The details of the interface between `GAMBIT`, `CLASS` and `plc` were presented in Ref. [31].

Crucially, we include the effects of ALP decays on the CMB power spectra. For $\tau_a > 10^{12}$ s, ALPs are still decaying when the electrons and protons in the primordial plasma combine to form neutral hydrogen. In this case the energy injected through these decays affects the free electron fraction and thereby delays recombination. The resulting constraints on decaying dark matter subcomponents were studied in e.g. Ref. [3]. The interfaces to the `CLASS` extension `ExoCLASS`, as well as to all modern versions of `CLASS` (v3.0 onwards) in `GAMBIT` provide a simple way to calculate these constraints for a wide range of models, based on transfer functions calculated with `DarkAges` [44]. The transfer functions used in this code, however, are designed to derive the impact of energy injection at redshifts around recombination and later, but not at the redshifts prior to recombination, at which spectral distortions primarily occur. At these redshifts, the universe is opaque and all injected energy will be transferred to heat up the plasma, such that it suffices to encode all impact of energy injection in a single effective efficiency factor f_{eff} . As we assume the ALP to decay into photons exclusively, this efficiency factor is assumed to be $f_{\text{eff}} = 1$. To include the constraints from energy injection in the present work we use the interface between `GAMBIT` and of `CLASS`, which was first developed in Ref. [31], further refined in Ref. [45], and which we updated to work with the latest version of `CLASS` (v3.1).

For $\tau_a < 10^{12}$ s the dominant effect of ALP decays is a change of the photon energy density relative to those of neutrinos and baryons, which will be discussed in more detail in the following.

3.1.1 Modifications of N_{eff} and η_b

In the absence of any beyond-the-SM physics, the only particle species that are still relativistic after BBN are photons and neutrinos. Hence the radiation energy density is given by

$$\rho_{\text{rad}}(T_\gamma, T_\nu) = \frac{\pi^2}{15} \left(T_\gamma^4 + 3 \times \frac{7}{8} T_\nu^4 \right), \quad (3.1)$$

where the first term accounts for the contribution from photons and the second term accounts for the contribution from neutrinos. As the two species are thermally decoupled, their respective temperatures T_γ and T_ν evolve independently. However, because all SM processes that could modify either of the two heat baths are inefficient after neutrino decoupling, the respective temperatures are only subject to redshifting. Their ratio is therefore fixed at neutrino decoupling. In the instantaneous-decoupling approximation, the ratio’s value can be derived by assuming conservation of entropy in both heat baths separately, such that $T_\nu^{\text{inst}} = (4/11)^{1/3} T_\gamma$. Non-instantaneous neutrino decoupling leads to a slightly larger temperature, and it is thus convenient to define the effective number of neutrino species via [46–49]

$$N_{\text{eff}} = \frac{3}{\left(\frac{4}{11}\right)^{4/3}} \left(\frac{T_\nu}{T_\gamma}\right)^4 = 3.0440 \pm 0.0002, \quad (3.2)$$

such that Eq. (3.1) can be rewritten as

$$\rho_{\text{rad}}(T_\gamma) = \frac{\pi^2}{15} \left[1 + \frac{7}{8} N_{\text{eff}} \left(\frac{4}{11} \right)^{\frac{4}{3}} \right] T_\gamma^4. \quad (3.3)$$

For ALPs with a lifetime $\tau_a > 10^4$ s, the value of N_{eff} during BBN will be equal to its SM prediction, given that the contribution of ALPs produced via the freeze-in mechanism to the total energy density during radiation domination is completely negligible. Subsequent ALP decays into photons will however increase T_γ relative to T_ν and thereby decrease the value of N_{eff} inferred from the CMB anisotropies. Furthermore, a modification of the photon temperature between BBN and recombination changes the baryon-to-photon ratio $\eta_b = n_b/n_\gamma$.

Including ALP decays, the evolution of the photon temperature is given by [7]

$$\frac{dT_\gamma}{dt} = \frac{15}{4\pi^2} \frac{m_a n_a(t) T_\gamma^{-3}}{\tau_a} - H(t) T_\gamma, \quad (3.4)$$

where $n_a(t)$ is the ALP number density, given by

$$n_a(t) = n_a(t_{\text{ini}}) \exp \left(-3 \int_{t_{\text{ini}}}^t dt' H(t') \right) e^{-t/\tau_a}, \quad (3.5)$$

where t_{ini} defines the initial time, which we take to be 10^3 s, i.e. sufficiently below the lower bound $\tau_a \geq 10^4$ s. The first term of Eq. (3.4) accounts for the temperature modification through the injection of entropy, whereas the second term accounts for the adiabatic cooling of the photon heat bath as the universe expands.

As $H(t)$ depends on both temperatures T_γ and T_ν , it is important to also track the evolution of the neutrino temperature T_ν as function of time. As ALPs do not interact with neutrinos, the temperature evolution is given by adiabatic cooling. Solving the corresponding differential equation yields

$$T_\nu(t) = T_\nu(t_{\text{ini}}) \exp \left(- \int_{t_{\text{ini}}}^t dt' H(t') \right). \quad (3.6)$$

We solve the evolution of the photon temperature in Eq. (3.4) numerically up to a final time $t_{\text{fin}} = 2 \times 10^{12}$ s, corresponding approximately to the time of matter-radiation equality.⁷ In Appendix B we provide details of our numerical implementation.

As the contribution from entropy injection is strictly positive, photons do not cool as fast as the neutrinos, and N_{eff} decreases between BBN and recombination. For $t_{\text{ini}} \leq t \leq t_{\text{fin}}$ the value of N_{eff} is

$$N_{\text{eff}}(t) = N_{\text{eff,BBN}} \left(\frac{T_\nu(t)}{T_\gamma(t)} \right)^4 \left(\frac{11}{4} \right)^{4/3}, \quad (3.7)$$

where $N_{\text{eff,BBN}} \equiv N_{\text{eff}}(t_{\text{ini}})$ is the value of N_{eff} at the end of BBN, given by Eq. (3.2). Since CMB observations are in agreement with the SM prediction, we can exclude parameter points that predict too large deviations from this value. To calculate the resulting constraints, we pass the Λ CDM parameters and the calculated value of $N_{\text{eff}}(t_{\text{fin}})$ to CLASS.

⁷Note that our calculation assumes that $\tau_a < t_{\text{fin}}$. For larger ALP lifetimes, the CMB will be directly sensitive to the exotic energy injection, and the resulting constraints are much stronger than those obtained from the modification of N_{eff} .

Furthermore, a change in T_γ also implies a change in the comoving photon number density and hence in the baryon-to-photon ratio. As $\eta_b(t_{\text{fin}}) \approx \eta_{b,0}$ is tightly constrained, we invert the calculation to obtain the baryon-to-photon ratio for earlier times $t_{\text{ini}} \leq t \leq t_{\text{fin}}$:

$$\eta(t) = \eta(t_{\text{fin}}) \left(\frac{T_\gamma(t_{\text{fin}})}{T_\gamma(t)} \right)^3 \exp \left(3 \int_t^{t_{\text{fin}}} dt' H(t') \right). \quad (3.8)$$

The baryon-to-photon ratio during BBN is then simply given by $\eta_{b,\text{BBN}} = \eta_b(t_{\text{ini}})$.

Note that the discussion above can be easily generalised to the case of additional contributions from dark radiation during BBN, viz. $N_{\text{eff,BBN}} = 3.045 + \Delta N_{\text{eff}}$. In this case the energy density in dark radiation scales exactly as the neutrino energy density, and the ratio $N_{\text{eff,BBN}}/N_{\text{eff,CMB}}$ is independent of the initial value of $N_{\text{eff,BBN}}$.

3.2 Primordial nucleosynthesis

Axion-like particles with $\tau_a > 10^4$ s decay after BBN, and the synthesis of primordial elements thus proceeds as in standard cosmology. Still, BBN constraints do play an important role for this model, as the decays of heavy ALPs with masses above $m_a \gtrsim 1$ MeV may contribute to the photodisintegration of some of the light elements [8]. The evolution of the primordial element abundances is thus affected by the injected EM spectrum through ALP decays and can be computed by solving the Boltzmann equation [50]

$$\frac{dY_N}{dt} = \sum_j \int_0^\infty dE f_\gamma (Y_j \sigma_{j\gamma \rightarrow N} - Y_N \sigma_{N\gamma \rightarrow j}), \quad (3.9)$$

where $N \in \{n, p, \text{D}, {}^3\text{H}, {}^3\text{He}, {}^4\text{He}, {}^6\text{Li}, {}^7\text{Li}, {}^7\text{Be}\}$ are the primordial elements, $Y_N = n_N/n_b$ their abundances, σ_p the cross sections for the various disintegration processes of the light elements, and f_γ the phase space distribution of nonthermal photons produced through ALP decays. Photons with energy $E_\gamma > m_e^2/(22T)$, where m_e denotes the electron mass, disappear efficiently via e^+e^- creation. Hence, disintegration processes only become efficient for temperatures $T < m_e^2/(22E_{\text{th}})$, where E_{th} denotes the threshold energy for each nuclear reaction. As typical threshold energies lie in the MeV range, photodisintegration becomes relevant only for $T \ll 10$ keV, well after the end of BBN. The various disintegration processes and their corresponding threshold energies are given in Table 1.

In scenarios where these depletion processes are sufficiently strong, the abundances of primordial elements will significantly differ from the predictions of vanilla Λ CDM, which are known to match observations in pristine gas clouds and metal-poor stars fairly well. Any sizeable modification of their abundances due to efficient disintegration will therefore be in considerable tension with measurements.

The element abundances are typically measured with respect to the hydrogen abundance (with the exception of ${}^3\text{He}$), and we thus express the relevant element abundances as $[\text{D}/\text{H}]$, $[{}^3\text{He}/\text{D}]$, $Y_p = [{}^4\text{He}/\text{H}]$ and $[{}^7\text{Li}/\text{H}]$.⁸

For most element abundances, the most up-to-date measurements can be found in [51], with the exception of $[{}^3\text{He}/\text{D}]$, for which we have computed our own improved estimation. To estimate $[{}^3\text{He}/\text{D}] = [{}^3\text{He}/\text{H}]/[\text{D}/\text{H}]$, we use the analysis of Ref. [52], who estimate the

⁸We do not consider the abundances of other primordial elements such as n , ${}^3\text{H}$ and ${}^7\text{Be}$ because they undergo further decay at late cosmological times and are totally depleted at the present time. We also do not consider ${}^6\text{Li}$, for which no reliable measurement of the primordial abundance exists.

Table 1: Threshold energies of dissociation processes considered in this work.

Reaction	E_{th} [MeV]	Reaction	E_{th} [MeV]
$\text{D} + \gamma \rightarrow n + p$	2.22	${}^6\text{Li} + \gamma \rightarrow {}^4\text{He} + n + p$	3.70
${}^3\text{H} + \gamma \rightarrow \text{D} + n$	6.26	${}^6\text{Li} + \gamma \rightarrow {}^3\text{H} + {}^3\text{He}$	15.79
${}^3\text{H} + \gamma \rightarrow 2n + p$	8.48	${}^7\text{Li} + \gamma \rightarrow {}^3\text{H} + {}^4\text{He}$	2.47
${}^3\text{He} + \gamma \rightarrow \text{D} + p$	5.49	${}^7\text{Li} + \gamma \rightarrow {}^6\text{Li} + n$	7.25
${}^3\text{He} + \gamma \rightarrow n + 2p$	7.12	${}^7\text{Li} + \gamma \rightarrow {}^4\text{He} + 2n + p$	10.95
${}^4\text{He} + \gamma \rightarrow {}^3\text{H} + p$	19.81	${}^7\text{Be} + \gamma \rightarrow {}^3\text{He} + {}^4\text{He}$	1.59
${}^4\text{He} + \gamma \rightarrow {}^3\text{He} + n$	20.58	${}^7\text{Be} + \gamma \rightarrow {}^6\text{Li} + p$	5.61
${}^4\text{He} + \gamma \rightarrow 2\text{D}$	23.84	${}^7\text{Be} + \gamma \rightarrow {}^4\text{He} + n + 2p$	9.30
${}^4\text{He} + \gamma \rightarrow \text{D} + n + p$	26.07		

protosolar abundance ratios based on their measures values today. To do so, we can use their estimate for protosolar $[{}^3\text{He}/\text{H}]$, and either an estimate for protosolar $[({}^3\text{He} + \text{D})/\text{H}]$ or a corrected direct measurement of Jovian $[\text{D}/\text{H}]$. Unlike Ref. [52], we use the weighted average of both methods, taking care when propagating the uncertainties. We find that $[{}^3\text{He}/\text{D}] = 0.82 \pm 0.11$. As the use of $[{}^3\text{He}/\text{D}]$ as a measurement is debated in the literature (see e.g. footnote 4 in Ref. [53]), we only use it as an upper limit by assigning an optimal likelihood to values smaller than the estimate. With this value and the recommendations from Ref. [51], the bounds that we use in this study are

$$\begin{aligned}
 [\text{D}/\text{H}] &= (2.547 \pm 0.025) \times 10^{-5} \\
 [{}^3\text{He}/\text{D}] &< 1.03 \text{ at } 95\% \text{ confidence level (CL)} \\
 Y_{\text{p}} &= 0.245 \pm 0.003 \\
 [{}^7\text{Li}/\text{H}] &= (1.6 \pm 0.3) \times 10^{-10} .
 \end{aligned} \tag{3.10}$$

3.2.1 Likelihood implementation and uncertainties

CosmoBit uses AlterBBN [54, 55] to compute an initial prediction of the primordial abundances of the light elements based on a given set of cosmological parameters. It then passes these abundances to ACROPOLIS [50]. The latter includes the effects of photodisintegration by computing the photon spectrum in the post-BBN thermal bath and solving the Boltzmann equation for each of the primordial element abundances. Details of the GAMBIT interface to ACROPOLIS can be found in Appendix A.3.

Both AlterBBN and our interface to ACROPOLIS compute ratios of abundances to the hydrogen abundance $[\text{Y}/\text{H}]$, from which we calculate $[{}^3\text{He}/\text{D}] = [{}^3\text{He}/\text{H}]/[\text{D}/\text{H}]$ in order to match the measured value of $[{}^3\text{He}/\text{D}]$ in Eq. (3.10). CosmoBit then computes a multi-dimensional Gaussian likelihood for all abundances, taking into account correlated uncertainties (see Ref. [31] for details):

$$-2 \ln \mathcal{L} = f(\mathcal{P}, \mathcal{O}) \mathcal{C}^{-1} f(\mathcal{P}, \mathcal{O})^T + (2\pi)^n \det(\mathcal{C}), \tag{3.11}$$

where \mathcal{P} is the vector of predicted abundances, \mathcal{O} the vector of observed abundances, and \mathcal{C} the covariance matrix. The function $f(\mathcal{P}, \mathcal{O})$ is introduced to differentiate elements for which the measurement is only an upper limit, such as ${}^3\text{He}$. Therefore, $f(\mathcal{P}, \mathcal{O})_i = \mathcal{P}_i - \mathcal{O}_i$

for all elements except for ${}^3\text{He}$, for which $f(\mathcal{P}, \mathcal{O})_{3\text{He}} = \max(\mathcal{P}_{3\text{He}} - \mathcal{O}_{3\text{He}}, 0)$. This way, the likelihood becomes flat when the predicted $[{}^3\text{He}/D]$ is smaller than the observed value, in line with our desire to implement the $[{}^3\text{He}/D]$ measurement as an upper bound only.

We construct the covariance matrix \mathcal{C} by computing the uncertainties for each of the abundances independently, and then combining them with correlation information. We estimate the uncertainties on each abundance for each parameter point by taking the larger of the distances between the mean and the high/low operational modes of `AlterBBN`, i.e. $\sigma_{Y_N} = \max(Y_N^{\text{high}} - Y_N^{\text{mean}}, Y_N^{\text{mean}} - Y_N^{\text{low}})$. A detailed computation of the correlations is possible in `AlterBBN`, but is computationally too expensive for our purposes. Fortunately, the production of primordial elements in our model proceeds as in ΛCDM . The correlations thus do not change significantly when varying the ΛCDM parameters within their typical ranges. We therefore employ a constant correlation matrix through our analysis that matches the one computed by `AlterBBN` with default settings. We then propagate the covariance matrix when computing the corrections after photodisintegration with `ACROPOLIS`, using the transfer matrix provided (see Appendix A.3 for details). Finally, in Eq. (3.11) we use a corrected covariance matrix, which matches the change to the observed ratio of abundances for $[{}^3\text{He}/D]$.

Besides their obvious use as predictions to compare with observations, the primordial abundances are also relevant for cosmological probes after BBN. In particular the primordial helium abundance Y_p affects the physics of recombination and thus is an input to the computation of the CMB power spectrum with `CLASS` (see Section 3.1). As recombination happens at temperatures much lower than the energy threshold for the photodisintegration reactions for Y_p , we make sure to include these effects on the predicted value of Y_p that we use for computing the spectrum of the CMB.

In addition to the BBN likelihood, we also include a nuisance likelihood for the neutron lifetime τ_n , as the reaction rates that produce the primordial elements are strongly affected by uncertainties in τ_n . In this analysis we use the laboratory measurement of τ_n from ultra-cold neutron (“bottle”) experiments, $\tau_{n,\text{bottle}} = (879.4 \pm 0.6) \text{ s}$, due to their smaller uncertainties. Effects of the uncertainty on τ_n on constraints from BBN, the CMB and other measurements were studied in Ref. [31].

3.2.2 The ${}^7\text{Li}$ problem

The BBN predictions for the abundance of ${}^7\text{Li}$ are known to be in disagreement with observations [56]. This is usually attributed to either systematic uncertainties in the determination of the primordial lithium abundance from astrophysical systems or uncertainties in the cross sections of nuclear reactions used in the theoretical calculations. Measurements of ${}^7\text{Li}$ are therefore generally not considered robust enough to be used for fits of cosmological data, and we do not include them in our scans unless explicitly stated otherwise.

Nevertheless, it is still interesting to consider physics beyond the SM as a possible solution of the ${}^7\text{Li}$ problem. For example, Ref. [8] suggests that the tension between predictions and observations can be resolved via ALP decays. The idea in this work is to consider ALPs in the mass range $3.17 \text{ MeV} < m_a < 4.44 \text{ MeV}$ such that the photons produced in their decay have enough energy for the photodisintegration of ${}^7\text{Be}$, but not enough for deuterium. We will return to this discussion in Section 4.

3.3 Spectral distortions

While the CMB radiation closely follows a black-body spectrum with characteristic temperature $T_{\text{CMB}} \approx 2.725 \text{ K}$ [57, 58], spectral distortions (SDs) exist due to interactions taking place

during the *entire* thermal history of the universe [4, 59–65]. As the relevance of the underlying physical processes that induce SDs changes over time, their redshift dependence has to be taken into account. This means that SDs are – at least in theory – sensitive to any known or new physics that alters the energy budget of the expanding universe, thus making SDs a useful, complementary probe of e.g. inflation, dark matter, or new particles (see Refs. [5, 66] for recent reviews).

Spectral distortions affect the photon spectrum $f(x, z)$, where $x = \omega/T_{\text{CMB}}$ is the dimensionless frequency as observed today and z is the redshift. This can be expressed as a deviation from the perfect black-body shape, i.e.

$$f(x, z) = \frac{1}{e^x + 1} + \Delta f(x, z). \quad (3.12)$$

Historically, SDs have been classified into different types, such as y -type distortions (from the Kompaneets or Sunyaev–Zeldovich parameter y) or chemical potential-type μ distortions. However, it is sufficient to specify the heating rates of the primordial plasma as a function of redshift, and integrate their spectral intensity $\Delta I \equiv (xT_{\text{CMB}})^3 \Delta f(x, z)/2\pi^2$ to compare to observations today. Given the frequency-integrated heating rate \dot{Q} from the Boltzmann equation, we have (see Ref. [5])

$$\Delta I \propto \frac{\Delta \rho_\gamma}{\rho_\gamma} = \int_0^\infty dz \frac{\dot{Q}}{(1+z)H\rho_\gamma}. \quad (3.13)$$

3.3.1 ALP distortion signal

Axion-like particles can induce SDs through ALP-photon conversion in the (known) Galactic magnetic field or a putative large-scale magnetic field of primordial origin [67–74], or through ALP decays [18, 19, 75]. The former process is only relevant for ALPs with $m_a \ll \text{eV}$. Since those ALPs are much lighter than the ones considered in this work, we only focus on the effect of decaying ALPs.

The effective energy injection rate of decaying ALPs, which contributes to the heating rate in Eq. (3.13), is given by [5]

$$\frac{dQ_{\text{inj},a}}{dt} = f_{\text{eff}} \xi \rho_{\text{DM}} \Gamma_a e^{-\Gamma_a t}, \quad (3.14)$$

where we set the efficiency factor $f_{\text{eff}} = 1$. This is a good assumption for decays before recombination [5]. In particular, $f_{\text{eff}} = 1$ is justified because the energy injection happens due to EM interactions, and we use the “on-the-spot approximation” of instantaneous energy injection at the time of the decay.

The ALP-induced SDs can in theory be calculated entirely via the unified approach of a full integration over the cosmological history outlined in Eq. (3.13). However, this is in practise a very expensive computation. The authors of Ref. [5] hence implemented a number of approximation schemes to simplify the calculation, which we briefly summarise in the next section. Examples for the ALP-induced spectral distortion signal in terms of ΔI are shown in the left panel of Fig. 3. There we choose ALP parameters that illustrate the power of using the general approach adopted in our analyses.

However, it is still useful to understand how ALPs can induce SDs by altering the single [76] and double [77] Compton rates. Compton scattering is efficient at redshifts $z \gtrsim 5 \times 10^4$, where μ -type distortions are sourced, and becomes inefficient at redshifts

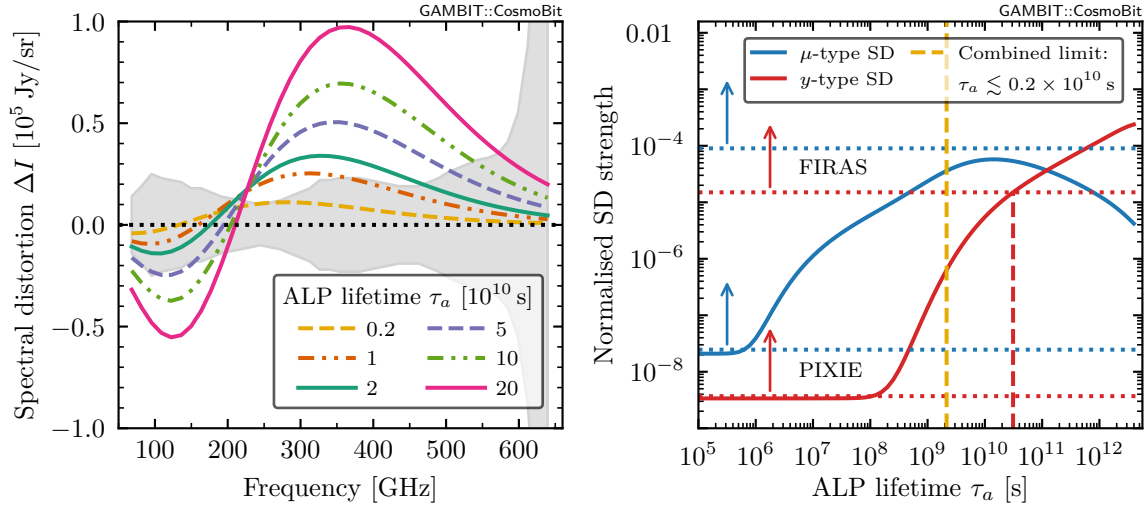


Figure 3: *Left:* Predicted SD (coloured lines) vs uncertainties of FIRAS observations (grey-shaded region). *Right:* Predicted μ and y distortions for ALPs with $m_a = 100$ MeV, $\xi = 5 \times 10^{-4}$, and $T_R = 5$ MeV as a function of τ_a . Horizontal dotted lines indicate the COBE/FIRAS limits and PIXIE forecasted limits; SDs above these lines are excluded at 95% CL. The vertical dashed lines indicate the limits from y -type SDs (red) and combined amount of SDs (gold), respectively.

$z \lesssim 5 \times 10^4$, leading to y -type distortions [78]. This approximate boundary corresponds to ALP lifetimes of $\tau_a \sim 10^{10}$ s, indicating that the type of ALP-induced SDs can change in different parts of the parameter space. Spectral distortions of a given type can thus only constrain the regions of parameter space where the underlying physical interactions are relevant.

If they apply, the ALP-induced μ [4] and y distortions [18] (see also Refs. [5, 19, 79]) can be approximated as

$$\mu \approx \frac{1}{\kappa_\mu} \left(\frac{3\rho_a}{\rho_\gamma} - \frac{8n_a}{n_\gamma} \right) \Big|_{T=T_{\text{dec}}}, \quad y \approx \frac{1}{4} \ln \left(\frac{\rho_a}{\rho_\gamma} + 1 \right) \Big|_{T=T_{\text{dec}}}, \quad (3.15)$$

with $\kappa_\mu \equiv 4\zeta(2)/\zeta(3) - 3\zeta(3)/\zeta(4) \approx 2.14$ and assuming that all ALPs decay at a temperature $H(T_{\text{dec}}) \sim \Gamma_a$.

If ALPs decay too fast or are too massive, they become inefficient in sourcing SDs, reducing the SDs in Eq. (3.15). The more accurately predicted ALP-induced μ and y distortions are shown in the right panel of Fig. 3. They should be compared to the constraints from COBE/FIRAS, which are $|\mu| < 0.9 \times 10^{-4}$ and $|y| < 1.5 \times 10^{-5}$ and at 95% CL [57, 80, 81]. We also included the expected limits for a null detection from a PIXIE-like experiment [5, Table 1].

While μ -type SDs typically set stronger constraints, this need not always be the case, as shown in Fig. 3. Here we chose ALP parameters where only y -type SDs result in a limit of $\tau_a < 3.1 \times 10^{10}$ s (vertical dashed red line). However, by deriving constraints using the total distortion ΔI (shown in the left panel of Fig. 3), we can even deduce a stronger limit of $\tau_a < 2.1 \times 10^9$ s (vertical dashed gold line). This illustrates the power of the spectral likelihood approach presented in the next section.

3.3.2 Likelihood implementation

The authors of Ref. [5] implemented the theoretical prediction of SDs as a part of CLASS v3.0 [41], incorporating adjusted versions of ExoCLASS [44], CosmoTherm [65, 82], and SZpack [83, 84]. They also provide a template likelihood for spectral distortions as a part of MontePython [85, 86], with specific realisations for COBE/FIRAS [57] data and PIXIE [87] forecasts.⁹

We pass the default options to MontePython for the SD likelihoods; in particular, this means that we use the ‘exact’ scheme for the SD branching approximation method. The ‘exact’ approximation scheme is based on the Green’s function method and a principle component analysis, which requires frequency binning [82, 88].

After specifying a frequency binning for each experiment, the likelihood for the observed distortions of photon spectral intensities $\Delta I_{\text{obs}}(\omega_i)$ at frequencies ω_i in the i th bin is

$$\ln \mathcal{L} = -\frac{1}{2} \sum_i \left(\frac{\Delta I_{\text{obs}}(\omega_i) - \Delta I_{\text{pred}}(\omega_i)}{\delta I(\omega_i)} \right)^2, \quad (3.16)$$

where the $\Delta I_{\text{pred}}(\omega_i)$ are the theory predictions by CLASS and the $\delta I(\omega_i)$ are the expected sensitivities, which are the sum in quadrature of the experimental noise level and theory uncertainties. This includes foreground maps, which were accounted for in Ref. [89] and largely based on Ref. [90].

One of the strengths of implementing the likelihood with the computational framework described above is that we avoid the complications of using Eq. (3.15) for μ and y distortions separately. Instead, we constrain all types of SDs simultaneously, using essentially all spectral information contained in the data. Our results should thus supersede previous works such as Refs. [19, 75].

3.4 ALP decays after SN1987A

The authors of Ref. [91] set limits on m_a and $g_{a\gamma}$ from ALPs emitted from core-collapse supernova SN1987A, using data [92] of the gamma-ray spectrometer [93] onboard the Solar Maximum Mission (SMM) [94]. The (instantaneous) emission spectrum of ALPs with energy E_a is

$$\frac{dN_a}{dE_a} = C \frac{E_a^2}{e^{E_a/T_{\text{SN}}} - 1} \sigma_{\text{P}}(m_a, g_{a\gamma}, E_a), \quad (3.17)$$

with $T_{\text{SN}} = 30.6 \text{ MeV}$ [91]. Here $C = 2.54 \times 10^{77} \text{ MeV}^{-1}$ is a fit to the results of Ref. [95], and σ_{P} is the Primakoff cross section [19], for which we use a Debye screening scale of $k_s = 16.8 \text{ MeV}$ [91]

The limits arise because the emitted ALPs can subsequently decay into two photons at some distance L_a , following an exponential distribution with mean distance

$$\ell = \frac{\beta\gamma}{\Gamma_a} = \frac{64\pi}{g_{a\gamma}^2 m_a^3} \sqrt{\frac{E_a^2}{m_a^2} - 1}, \quad (3.18)$$

where β and γ are the dimensionless ALP velocity and Lorentz boost factor, respectively.

⁹The PIXIE likelihood is based on the instrument’s expected sensitivity and sets the data equal to the expected signal without any new, exotic physics (a null detection). For FIRAS, the authors of Ref. [5] use the observed sensitivity but effectively also assumes a null detection. We expect this approximation to be reasonable as the observed residuals are small compared to the uncertainties shown in the left panel of Fig. 3.

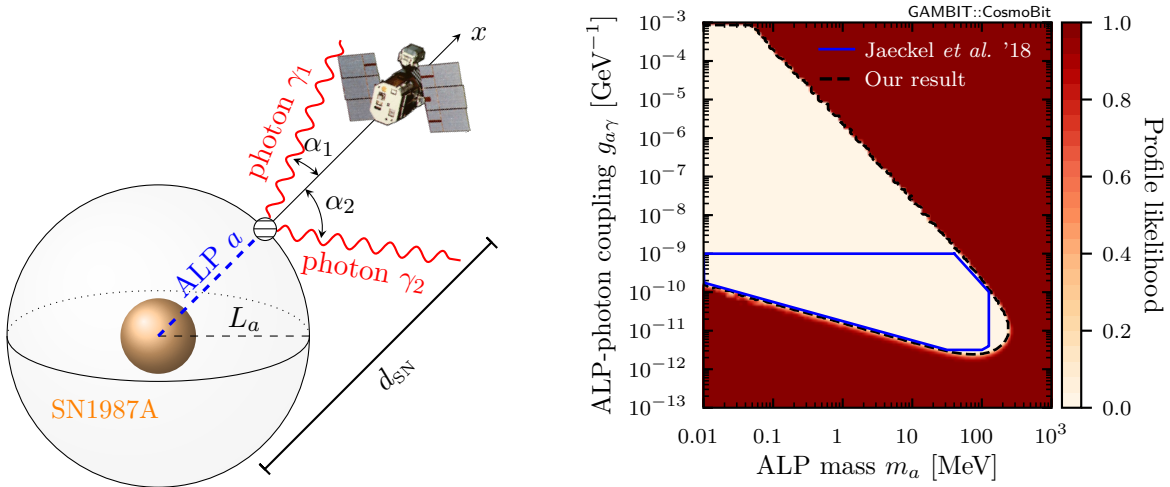


Figure 4: *Left:* Geometry for SN1987A ALPs decaying into two photons, emitted at angles $\alpha_{1/2}$ at a distance L_a in the reference frame of the SMM satellite. Image credit for the SSM satellite to G. Nelson/NASA (JSC image library; public domain). *Right:* Our profile likelihood and exclusion region of m_a and $g_{a\gamma}$ (dashed black line) compared to that of Ref. [91] (blue line) at 3σ CL, assuming 1 degree of freedom.

To calculate the expected photon fluence from ALP decays, we estimate the “success rate”, i.e. the fraction of photons arriving at the SMM satellite, via Monte Carlo simulations. The expected fluence is then calculated by integrating Eq. (3.17), dividing by the surface of the sphere centred on the supernova with the distance to the detector as the radius ($d_{\text{SN}} = 51.4$ kpc), and multiplying by the success rate. An illustration of the setup is provided in the left panel of Fig. 4.

Following Ref. [91], we use a Gaussian likelihood function assuming a null result where the standard error is derived from the observed photon fluence in the energy range of 25–100 MeV. The resulting exclusion limit is shown in the right panel of Fig. 4. We provide details of our implementation in Appendix C, as well as a table of the signal prediction that ships as part of the latest version of GAMBIT, and publicly available Python scripts for our Monte Carlo procedure.¹⁰

3.5 Stellar evolution and cluster counts

Due to their weak interactions with matter, ALPs can carry away energy from stellar interiors, affecting the evolution of e.g. Horizontal Branch (HB) stars by reducing their lifetimes [96–98]. This allows one to place a limit on $g_{a\gamma}$ by comparing the number of HB stars N_{HB} to the number of stars on the Red Giant Branch (RGB) N_{RGB} in globular clusters (GCs). The associated “ R parameter” observable $R_{\text{GC}} \equiv N_{\text{HB}}/N_{\text{RGB}}$ [99], is directly connected to the lifetimes and energy loss rates of these systems.

The theoretically expected R parameter can be estimated from stellar evolution simulations that include the energy loss from ALPs [23, 34, 99, 100],

$$R_{\text{GC}} \approx 0.022 - 0.443 (1 + 0.965 |g_{a\gamma}|)^{1/2} + 7.331 Y_{\text{GC}}, \quad (3.19)$$

¹⁰Available on Github at <https://github.com/marie-lecroq/ALP-fluence-calculation>.

where $Y_{\text{GC}} = 0.2515(17)$ [101] is an estimate for the Helium abundance in GCs. The observed value of the R parameter, $R_{\text{GC,obs}} = 1.39 \pm 0.03$ [34], is consistent with the axion-free value from (3.19), which is $R_{\text{GC}} = 1.42 \pm 0.01$.

Equation (3.19) is only valid for ALP masses much smaller than the typical core temperature of an HB star, $m_a \ll T_{\text{HB}} \sim 10 \text{ keV}$. For greater ALP masses, the Primakoff interaction cross section, which leads to the energy loss in HB stars, should be multiplied by a correction factor $F(m_a)$. Using Eq. (A.5) of Ref. [19], we obtain $F(m_a)$ normalised such that $F(m_a \ll T_{\text{HB}}) \simeq 1$. As the Primakoff cross section and the energy loss are proportional to $g_{a\gamma}^2$, it is sufficient to replace $g_{a\gamma} \mapsto g_{a\gamma} \times \sqrt{F(m_a)}$ in Eq. (3.19) in order to apply the correction.

The R parameter likelihood is part of the DarkBit [102] module of GAMBIT. It was first introduced in Ref. [103], and the correction for larger masses was added in Ref. [104]. Also note that slightly stronger limits in the region around $m_a \sim 100 \text{ keV}$ may be obtained by considering photon coalescence [105]. As the R parameter constraints only marginally affect our results, we do not include these slightly improved constraints.

3.6 Other likelihoods

Finally, we make use of MontePython v3.5 to include the BOSS DR12 [106] likelihood for Baryon Acoustic Oscillations and the Pantheon [107] likelihood for type Ia supernovae. The latter likelihood requires the introduction of an additional nuisance parameter M corresponding to the absolute magnitude of a type Ia supernova, see Ref. [31] for details.

4 Results

By combining all of the likelihoods introduced in Section 3, we are now in the position to perform a global scan of ALP decays in the early universe. We first present the results from a frequentist analysis of the parameter space, which enables us to identify particularly interesting parameter regions, and then perform a Bayesian analysis, which provides complementary information on the preference between the ALP model and ΛCDM .

4.1 Frequentist scans

The parameters included in our scans are summarised in Table 2. We emphasise that the parameter ranges for the ALP model are somewhat arbitrary and are chosen in such a way as to capture all relevant features.¹¹ In other words, while viable parameter regions exist also beyond the scan ranges, they do not lead to any qualitative changes in our conclusions. The scan ranges for the ΛCDM and other nuisance parameters are chosen so that they fully encompass the allowed parameter regions.

To perform comprehensive scans of the parameter space we employ the differential evolution sampler Diver v1.0.4 [108], using a population size of 20,000 and a convergence threshold of 10^{-8} . We carried out additional scans targeted at specific regions of parameter space in order to ensure that the likelihood function is fully explored. As we perform a frequentist analysis based on the profile likelihood, we simply combine samples from different

¹¹We note that for the smallest ALP masses and the shortest lifetimes that we consider, the assumption of non-relativistic ALP decays is not a good approximation for ALPs produced via freeze-in, because it underestimates the energy density of ALPs. Nevertheless, we find that this parameter region is robustly excluded in our scans. A more accurate treatment would only strengthen this exclusion and is therefore unnecessary.

Table 2: List of model and nuisance parameters included in our scans and their corresponding scan ranges.

Model parameter		Scan range	
ALP mass	m_a	[0.001, 200]	MeV
ALP lifetime	τ_a	$[10^4, 10^{13}]$	s
ALP abundance	ξ	$[10^{-12}, 10^2]$	
Baryon abundance	ω_b	[0.020, 0.024]	
Dark matter abundance	ω_{DM}	[0.10, 0.13]	
Hubble constant	H_0	[62, 74]	$\text{km s}^{-1} \text{Mpc}^{-1}$
Redshift of reionisation	z_{reio}	[4.5, 9.5]	
Primordial curvature	$\ln(10^{10} A_s)$	[2.9, 3.2]	
Scalar spectral index	n_s	[0.9, 1.1]	
Neutron lifetime	τ_n	[875, 895]	s
Planck nuisance parameter	A_{Planck}	[0.9, 1.1]	
Pantheon nuisance parameter	M	[-20, -18]	

scans. The results shown below are based on a total of 33.5 million samples and took over 700,000 CPU hours on the Joliot-Curie HPC cluster at CEA (France).

In what follows, we use Wilks’ theorem to determine the exclusion regions of the profile likelihood. To do so we assume that the number of degrees of freedom corresponds to the number of free parameters – i.e. two degrees of freedom for the two-dimensional plots shown below. It is well known that the assumptions of Wilks’ theorem often do not hold exactly, and that care should be taken to ensure its validity when claiming a discovery [e.g. 109]. For the purpose of setting limits w.r.t the best-fitting point, however, we argue that Wilks’ theorem typically presents a good approximation and avoids costly Monte Carlo simulations that do not alter significantly the findings of our analysis (see e.g. Ref. [104, appendix D] for a demonstration in the context of global fits).

Our main results are shown in Fig. 5, where we plot the profile likelihood as a function of different combinations of ALP model parameters. We remind the reader that there are three independent ALP model parameters, one of which is profiled out in each panel. We use grey shading to indicate parameter regions that are not explored in our scans due to the chosen scan ranges for m_a and τ_a . In particular, the upper bound on τ_a implies a lower bound on $g_{a\gamma}$ as a function of m_a ; this, in turn, implies a lower bound on ξ because of the irreducible freeze-in contribution, cf. Eq. (2.8).

Each of the applied constraints is sensitive to a specific combination of model parameters and can therefore be understood most naturally in one of the panels in Fig. 5. For example, the astrophysical constraints are independent of ξ and therefore directly constrain the combination $(m_a, g_{a\gamma})$ or, equivalently, $(\tau_a, g_{a\gamma})$. The cosmological constraints, on the other hand, are to first approximation independent of m_a . They are therefore most easily expressed in terms of (ξ, τ_a) . Coloured lines and labels indicate the locations of individual constraints (at 99.7%

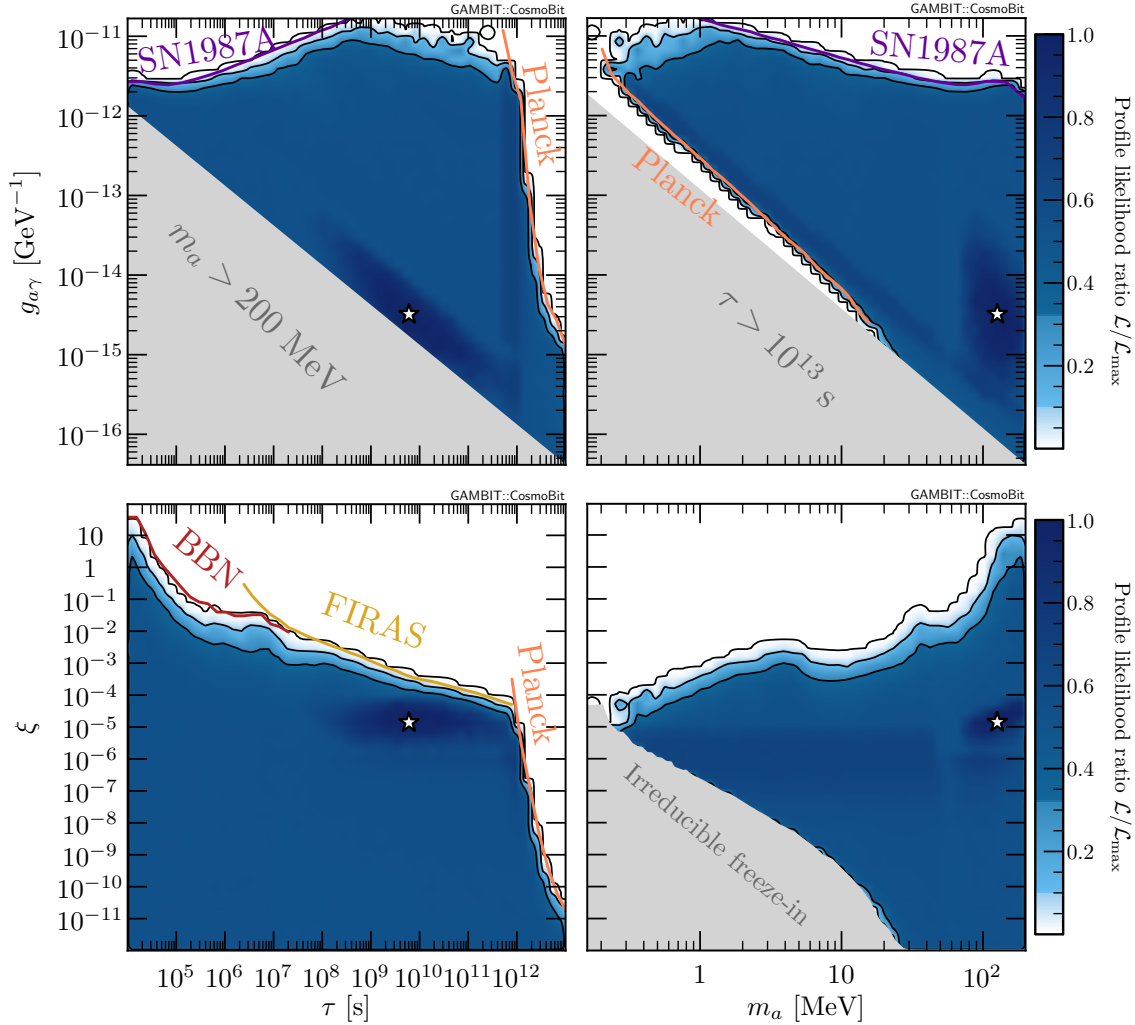


Figure 5: Profile likelihood for various combinations of the parameters. Black lines correspond to the 1σ , 2σ , and 3σ contours around the best-fitting point (white star), whereas coloured lines are 99.7% CL limits from the individual likelihoods. Shaded grey regions are excluded from our study as they are beyond the scan ranges.

CL) in the various panels.

We find that combining the various likelihoods excludes large parts of the ALP parameter space. The cosmological constraints become stronger with increasing ALP lifetime and require $\xi \ll 1$, unless $m_a \gtrsim 100$ MeV and $\tau_a \lesssim 10^5$ s. Moreover, we find that there are no viable parameter regions for $m_a \lesssim 300$ keV as a result of the SN1987A constraint. This is because it places an upper bound on $g_{a\gamma}$, and therefore a lower bound on τ_a , which in turn leads to strong cosmological constraints. We emphasise that this lower bound on m_a is a direct result of observational constraints and does not depend on the scan ranges that we have chosen. However, when significantly expanding the scan ranges, we find that there is a second disjoint region of allowed parameter space at smaller ALP masses, ALP lifetimes comparable to or greater than the age of the universe, and a tiny abundance. The most promising way to probe this parameter region would be through searches for X-ray lines [e.g. 19], which we do not

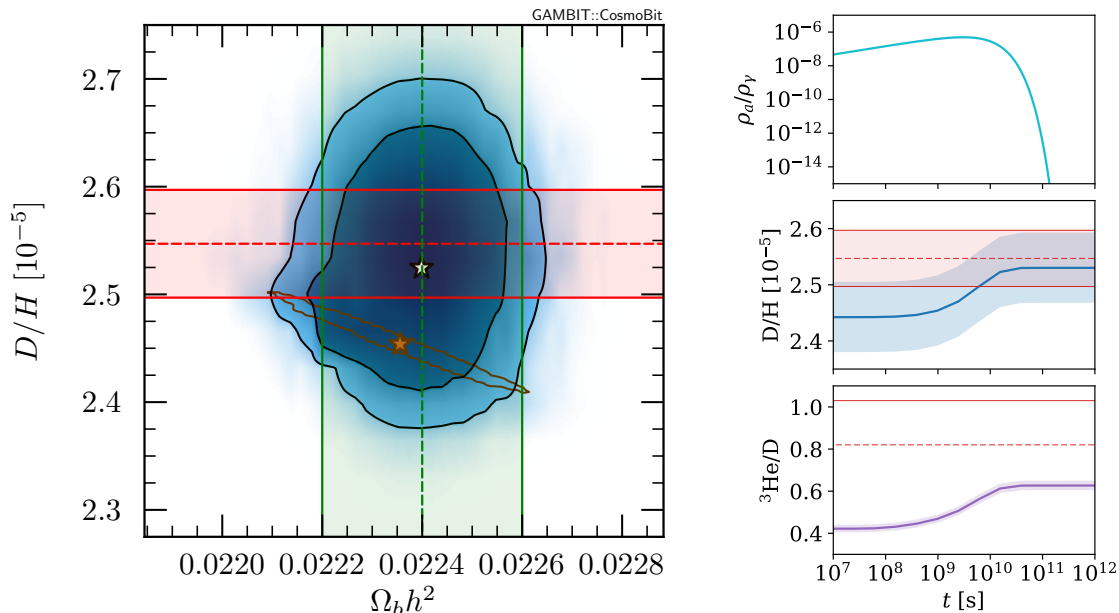


Figure 6: *Left:* Profile likelihood in the plane of baryon density $\Omega_b h^2$ and deuterium abundance D/H . Black lines correspond to the 1 σ and 2 σ contours around the best-fitting point for the ALP model (white star). The brown line indicates the 2 σ contour around the best-fitting point in Λ CDM (brown star). Horizontal and vertical lines mark the central values and the 2 σ intervals of the measured $\Omega_b h^2$ [42] and D/H (see Section 3.2) values, respectively. *Right:* Time evolution of the ALP energy density (relative to the energy density of photons) as well as of the abundance ratios D/H and $^3\text{He}/\text{D}$ for the best-fit point. Dashed (solid) red lines indicate the observed values (2 σ bounds) of the abundance ratios, the shaded bands represent theoretical uncertainties, which for the case of D/H are comparable to the measurement uncertainties.

consider in the present work.

We observe that it is essential to combine different likelihoods in order to constrain the parameter space from various directions. However, we find that no relevant constraints are obtained from the modification of N_{eff} and η_b discussed in Section 3.1.1. The reason is that in the parameter regions allowed by all other constraints, the energy density of ALPs relative to radiation is simply too small to change the photon temperature in an observable way. Even for the shortest ALP lifetimes and largest ALP masses that we consider, eq. (2.9) implies $\rho_a/\rho_{\text{tot}} < 10^{-3}$, which restricts the relative modification of N_{eff} and η_b to a similar level. Furthermore, we find that once all other constraints are imposed, the likelihood for R_{GC} has no effect on the remaining parameter space, since in the allowed parameter region the ALPs are too heavy to be produced in horizontal branch stars.

We find that large regions of parameter space feature a relatively constant likelihood, which corresponds to the case that the ALP has negligible impact on cosmological observables. However, we also find a small region that is slightly preferred by data. The best-fitting point is found at $m_a = 126.1$ MeV, $\tau_a = 6.04 \times 10^9$ s (corresponding to $g_{a\gamma} = 3.31 \times 10^{-15}$ GeV $^{-1}$), and $\xi = 1.48 \times 10^{-5}$ (compatible with the irreducible freeze-in contribution for $T_{\text{R}} \lesssim 4$ GeV) and gives $-2\Delta \ln \mathcal{L} = 1.51$ compared to the case of negligible ALP abundance.

To understand the origin of this improvement, we show in Fig. 6 the allowed parameter region in terms of the deuterium abundance D/H and the baryon abundance $\Omega_b h^2$. For

comparison, we show the measured 2σ CL interval for both quantities as horizontal and vertical bands (see Section 3.2 for the experimental measurement of D/H and Ref. [42] for that of $\Omega_b h^2$). On the one hand, we find that in Λ CDM these two quantities are tightly correlated, and it is not possible to perfectly fit both measurements simultaneously. On the other hand, in our ALP model it is possible to increase the deuterium abundance through photodisintegration of ^4He , which clearly improves the agreement with observations.

We remind the reader that our scans include a theoretical uncertainty for the prediction of D/H, which is calculated on the fly using `AlterBBN`. Indeed, for most points the relative theoretical uncertainty is large compared to the observational uncertainty (see also Ref. [31]). This is why the allowed parameter region of our model extends well beyond the observed confidence interval for D/H. The best-fitting point of Λ CDM is hence found to lie within the 1σ allowed parameter region of the ALP model. This finding emphasises the need for an accurate treatment of uncertainties in the calculation of BBN abundances. Indeed, neglecting the theoretical uncertainty on D/H would significantly overestimate the preference of the ALP model over Λ CDM.

While the preference for our model over Λ CDM is certainly not significant, it is still interesting to understand which future observations may shed further light on this issue. We therefore performed another set of scans where we replace the COBE/FIRAS likelihood with a mock likelihood corresponding to the sensitivity of a PIXIE-like mission [87]. Such a PIXIE mock likelihood, where a null result for new physics was assumed to generate typical observed data, was first studied in Ref. [5] and made publicly available through `MontePython`.¹² The results are shown in Fig. 7 and are indicative of the improvements that can be expected from future measurements of CMB SDs. We find that a null result from PIXIE would indeed significantly tighten the constraints on our model and robustly exclude the best-fit point found in the previous scan. In other words, the best-fitting point of our main analysis including the COBE/FIRAS likelihood predicts CMB SDs that should be easily detectable with future missions, such that PIXIE would be able to determine to high significance whether photodisintegration of ^4He caused by ALP decays is responsible for the large value of D/H compared to the naive Λ CDM prediction. We emphasise that future measurements of CMB anisotropies (using for example CMB-S4 [110, 111]) would not lead to similar improvements in the constraints, since the predicted modifications of ΔN_{eff} and η_b are simply too small to be observable.

Let us finally comment on the ^7Li problem. To understand if the ALP model considered in this work can solve this problem, we ran a separate scan in which the ^7Li measurement is included as an additional likelihood. In agreement with the results from Ref. [8], we find that the ALP abundances required to significantly modify the ^7Li abundance must be of order $\xi \sim 1$. Such large ALP abundances are however robustly excluded by CMB SDs, which were not considered in Ref. [8]. We therefore conclude that, once all relevant constraints are taken into consideration, ALPs decaying into photons cannot resolve the ^7Li problem.

To conclude, we note that we have checked explicitly that all results presented in this section are robust against additional contributions to ΔN_{eff} from further ultrarelativistic species. This is a direct consequence of the fact that the constraint on N_{eff} does not dominate

¹²Note that it is not currently possible in `MontePython` to combine the COBE/FIRAS likelihood and the PIXIE mock likelihood. However, since the latter is assumed to be substantially more constraining than the former, we expect results from a combination of both likelihoods to be very similar to the case with only the PIXIE mock likelihood. In any case, the PIXIE likelihood is subject to considerable unknowns at this point and should therefore only be considered an approximation.

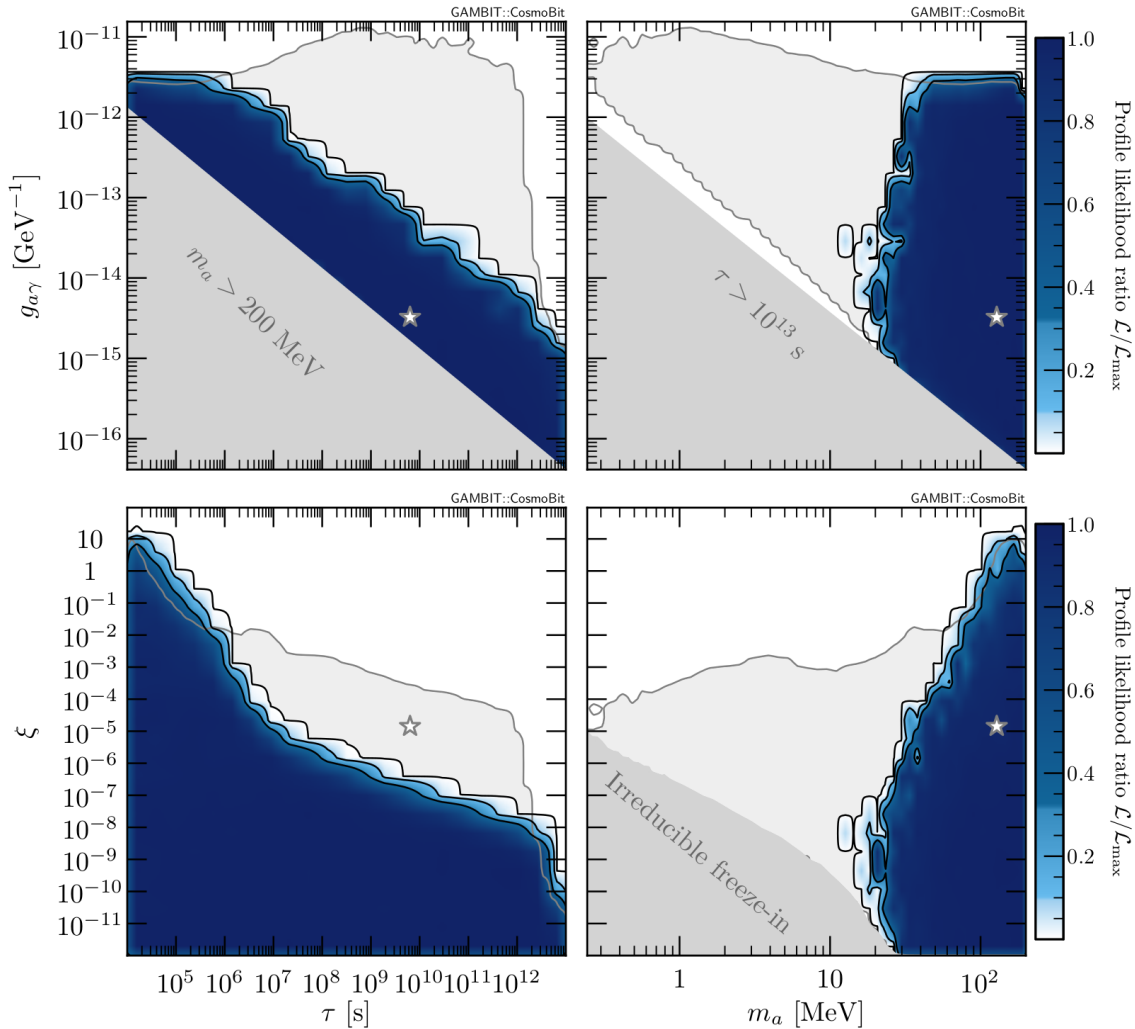


Figure 7: Profile likelihoods as in Fig. 5 but for the projected PIXIE limits instead of the FIRAS likelihood. Black lines again correspond to the 1σ , 2σ , and 3σ contours around the best-fit point. The likelihood in the allowed region is very flat such that the best-fit point is not unique; it is thus not shown here. The grey shaded region corresponds to the 2σ region around the best-fit point (grey star) of the scans with the FIRAS likelihood from Fig. 5.

our exclusion limits.

4.2 Bayesian scans

The priors that we employ for Bayesian scans extend across the ranges quoted in Table 2. For the Λ CDM parameters and the nuisance parameters, all of which are relatively tightly constrained by data, we employ flat priors on each interval.¹³ The ALP model parameters, on the other hand, can vary over many orders of magnitude. In order not to impose a specific scale, we employ logarithmic priors for these parameters.

¹³In particular, we use a flat prior for A_{Planck} and include a Gaussian likelihood in our scans to constrain it. This procedure leads to the same posteriors as when using a Gaussian prior with no additional likelihood.

These scans took place on the Cambridge CSD3 DiRAC cluster, with sampling by PolyChord [112] v1.20.1 with `nlive= 1000`, `nprior= 40000`, `tol= 10-10` and `logzero= -10101`. These settings ensure good sampling of the prior, the likelihood peak and the broader posterior, and give control over the points thrown away at the prior level.

We performed a Λ CDM run and an ALP run with the same priors on the cosmological parameters. The resulting posteriors from the two runs for the baseline cosmological parameters are indistinguishable and no correlations are found between the Λ CDM parameters and the ALP parameters. For the ALP run, we consider two effective priors on the ALP parameters: the default “box prior” given by the ranges in Table 2, and a “Primakoff prior”, which results from throwing away points within the box prior where the irreducible ALP abundance from the Primakoff process (for $T_R = 5$ MeV) is incompatible with the chosen value of ξ . We emphasise that in the case of the “box prior” we still impose the lower bound on ξ at the level of the likelihood, so that the two approaches lead to indistinguishable posteriors.

Fig. 8 shows the Bayesian equivalent of Fig. 5. The key difference here compared to Fig. 5 is that both the priors and the posterior are plotted in Fig. 8. Note that whilst these panels do show individual points, they do not show all points sampled. The points shown are an equal number of equal-weight draws from each distribution. This is in contrast to Fig. 5, which is based on profiling the likelihoods of all sample points within fixed-width bins, and interpolating the resulting maximum likelihood between them.

For a likelihood \mathcal{L} and prior π , one may compute a Bayesian evidence (or marginal likelihood) \mathcal{Z} , posterior distribution \mathcal{P} from Bayes theorem, and Kullback–Leibler divergence \mathcal{D}_{KL} between the prior and posterior distribution [113]:

$$\mathcal{P} = \frac{\mathcal{L}\pi}{\mathcal{Z}} \quad \mathcal{Z} = \int \mathcal{L}\pi d\theta = \langle \mathcal{L} \rangle_{\pi} \quad \mathcal{D}_{\text{KL}} = \int \mathcal{P} \ln \frac{\mathcal{P}}{\pi} d\theta = \left\langle \ln \frac{\mathcal{P}}{\pi} \right\rangle_{\mathcal{P}} \quad (4.1)$$

In order to quantitatively compare the two models, Λ CDM and Λ CDM +ALPs, we compute the Bayesian evidence, \mathcal{Z} for the different prior definitions:

$$\begin{aligned} \ln \mathcal{Z}_{\text{ALPs}}^{\text{box}} &= -1039.09 \pm 0.14 \\ \ln \mathcal{Z}_{\text{ALPs}}^{\text{Primakoff}} &= -1038.12 \pm 0.14 \\ \ln \mathcal{Z}_{\Lambda\text{CDM}} &= -1037.11 \pm 0.14 \end{aligned} \quad (4.2)$$

Λ CDM is thus marginally preferred to the ALP model in a Bayesian sense by $\Delta \ln \mathcal{Z}^{\text{box}} = -1.98 \pm 0.20$ or $\Delta \ln \mathcal{Z}^{\text{Primakoff}} = -1.02 \pm 0.20$ log units. The error bars in this section all arise from the nested sampling procedure [114, 115] as a result of the probabilistic integration used for estimating the evidence.

We can see where the difference in Bayesian evidence arises by using the Occam’s razor equation to decompose the evidence into a goodness of fit (average log-likelihood) and Occam penalty (Kullback–Leibler divergence) [116]. By taking logarithms and rearranging Bayes theorem before posterior averaging, we recover the Occams razor equation, which in our case decomposes the evidences as

$$\begin{aligned} \ln \mathcal{Z} &= \langle \ln \mathcal{L} \rangle_{\mathcal{P}} - \mathcal{D}_{\text{KL}} \\ \ln \mathcal{Z}_{\text{ALPs}}^{\text{box}} &= -1012.27 - 26.82 \\ \ln \mathcal{Z}_{\text{ALPs}}^{\text{Primakoff}} &= -1012.27 - 25.85 \\ \ln \mathcal{Z}_{\Lambda\text{CDM}} &= -1012.38 - 24.72 \end{aligned} \quad (4.3)$$

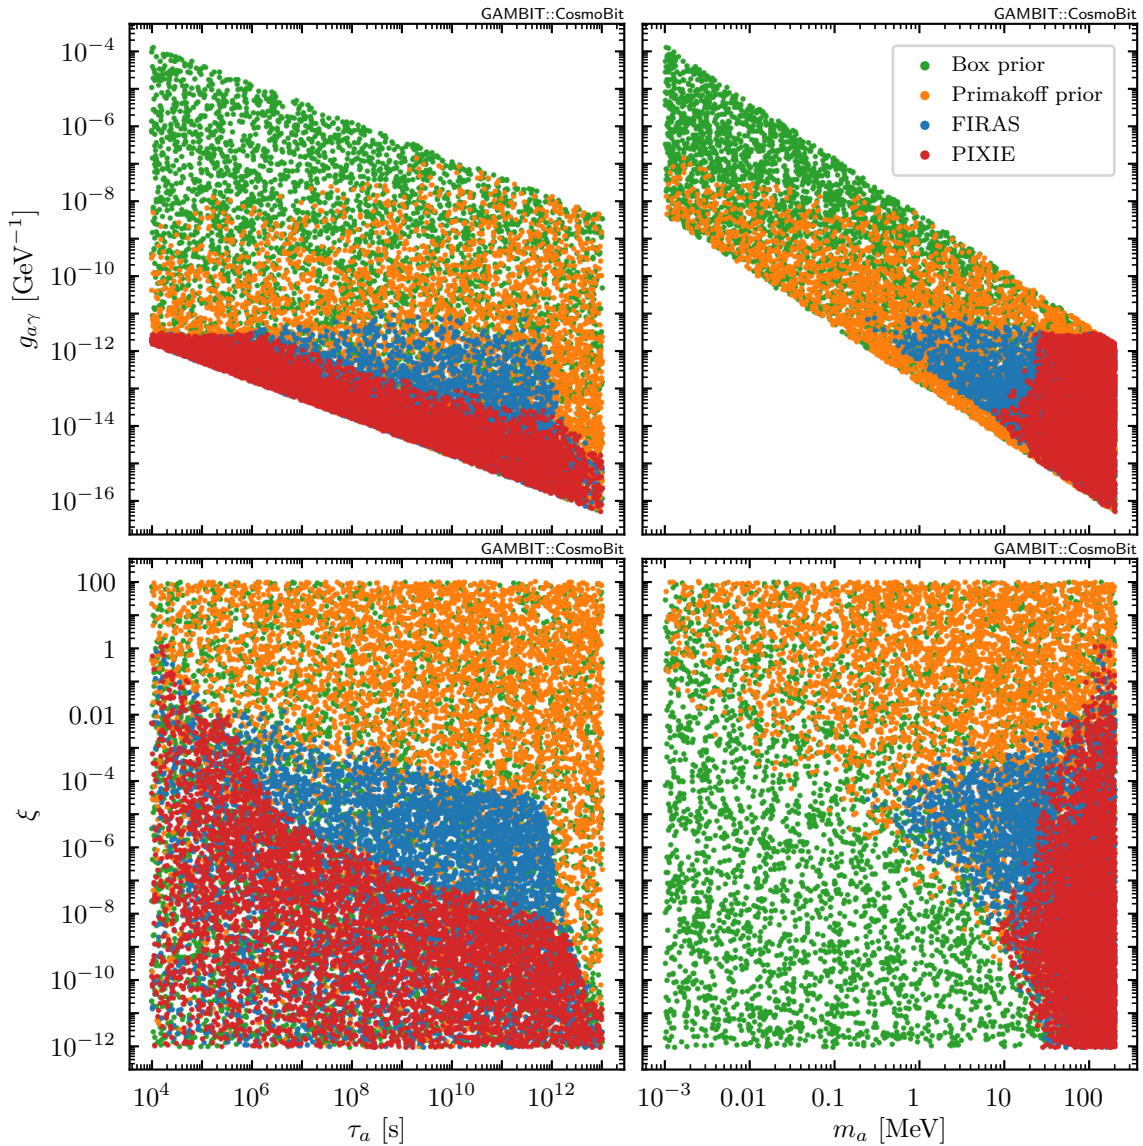


Figure 8: Bayesian scan analogue of Fig. 5. Green points represent equally-weighted samples drawn from the default box prior, orange points are samples drawn from the more restrictive Primakoff prior (see text for details), and blue and red points are samples drawn from the posterior using the COBE/FIRAS and PIXIE-type likelihoods, respectively.

These results demonstrate that all models in a Bayesian sense fit the data equally well, with $\Delta \langle \ln \mathcal{L} \rangle_{\mathcal{P}} = 0.11 \pm 0.07$ relative to Λ CDM, consistent with zero. The ALPs are penalised by having additional constrained parameters, with an additional compression of $\Delta \mathcal{D}_{\text{KL}}^{\text{box}} = 2.09 \pm 0.19$ or $\Delta \mathcal{D}_{\text{KL}}^{\text{Primakoff}} = 1.12 \pm 0.19$, i.e. the parameter space contained in the posterior is a factor of $\exp(\Delta \mathcal{D}_{\text{KL}}^{\text{box}}) = 8.3 \pm 1.5$ or $\exp(\Delta \mathcal{D}_{\text{KL}}^{\text{Primakoff}}) = 3.1 \pm 0.6$ smaller than the prior. This compression factor measures how well the data constrain the model. Note that employing the more stringent Primakoff prior rules out a little more than half the parameter space in advance, and therefore reduces the Occam penalty. If the COBE/FIRAS data set is replaced with the PIXIE forecast, the only change is an increase in parameter space

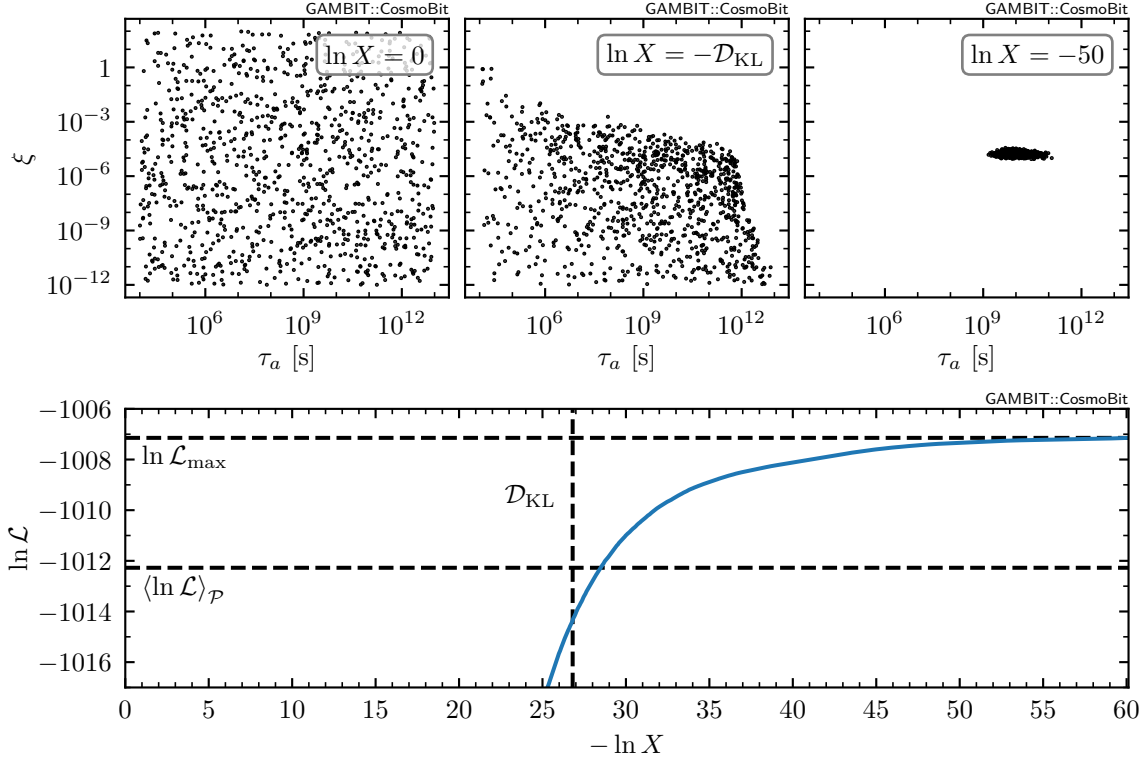


Figure 9: *Top:* Live point locations in the (τ_a, ξ) plane during different stages of nested sampling (NS). The i th NS iteration may be parametrised by the log-prior volume $\ln X \approx -i/n_{\text{live}}$. Three critical stages are shown: the initial (non-Primakoff) prior points when $\ln X = 0$, at the “typical set” where $\ln X = -\mathcal{D}_{\text{KL}}$, and at a much lower value around the point where NS terminates according to our selected tolerance (which in this case is around the best-fit point). The prior points have large volume and low likelihood, whereas the peak has high likelihood and low volume. The Bayesian “sweet spot” (i.e. majority contribution to the evidence) occurs around the typical set with approximate parameter volume given by $e^{\mathcal{D}_{\text{KL}}}$. The higher-likelihood phase therefore occupies a parameter region $\approx 10^7$ times smaller than the usual posterior, so is not preferred in a Bayesian sense and does not appear in Fig. 8. *Bottom:* Phase plot of log-likelihood against log-volume, showing the average log-likelihood $\langle \ln \mathcal{L} \rangle_{\mathcal{P}}$ over the posterior \mathcal{P} , the Kullback–Leibler divergence \mathcal{D}_{KL} , and the maximum log-likelihood value $\ln \mathcal{L}_{\text{max}}$. The high likelihood peak can be identified as beyond the region where there is a small inflection point at $\ln X \approx -42$. All of this information is extractable in post-processing with the NS Python package `anesthetic` [117].

compression by approximately a factor of e : $\Delta \mathcal{D}_{\text{KL}}^{\text{box}} = 3.08 \pm 0.19$ and $\Delta \mathcal{D}_{\text{KL}}^{\text{Primakoff}} = 2.11 \pm 0.19$, i.e. the posterior is a factor of $\exp(\Delta \mathcal{D}_{\text{KL}}^{\text{box}}) = 22.2 \pm 4.2$ or $\exp(\Delta \mathcal{D}_{\text{KL}}^{\text{Primakoff}}) = 8.4 \pm 1.6$ smaller than the prior. The Likelihood contribution to the evidence remains the same.

Fig. 9 shows that using the full scope of the nested sampling run, we can gain insight into the high likelihood region X shown in Fig. 5. This plot showcases the Bayesian interplay between likelihood and parameter volume.

5 Conclusions

Axion-like particles (ALPs) are well-motivated candidates for light new physics and arise generically in many extensions of the Standard Model. In the present work we studied ALPs

with masses in the keV–MeV range. We focused on ALPs with couplings to photons small enough to be inaccessible to accelerator experiments, but large enough to ensure that the ALPs decay before recombination and therefore do not contribute to the observed dark matter abundance. In this coupling range, ALPs can affect various cosmological and astrophysical observables, which makes it possible to search for them and to constrain their properties.

A central aspect of our analysis is that we did not assume that ALPs enter into thermal equilibrium in the early universe, such that their abundance is essentially an independent parameter. Nevertheless, one can place a lower bound on the abundance from the irreducible freeze-in production, and we provided a detailed calculation of this abundance as a function of the ALP mass and reheating temperature. At the same time, BBN and CMB constraints (both from anisotropies and spectral distortions) place tight upper bounds on the ALP abundance, typically requiring that these particles constitute only a tiny fraction of the dark matter abundance even before they decay.

To perform a global analysis of the ALP parameter space, we interfaced the GAMBIT framework with a number of external codes in order to combine state-of-the-art likelihood calculations for all relevant data sets. Most notably, for the first time we combine detailed calculations of the effects of photodisintegration of light elements with constraints on CMB spectral distortions, both of which give highly relevant constraints on ALP models. Our setup makes it possible to scan simultaneously over the ALP model parameters, the Λ CDM parameters, and various nuisance parameters, in order to identify the most interesting regions of parameter space (see Fig. 5).

In particular, the decays of ALPs into photons can modify both the abundances of light elements and the baryon-to-photon ratio η_b . As a result, we find that $[D/H]$ and $\Omega_b h^2$ are less tightly correlated in our model than in Λ CDM (see Fig. 6). This makes it possible to achieve a slightly better fit to observations in our model than in Λ CDM, although the Bayesian evidence does not prefer the ALP model over standard Λ CDM. However, we did not find any allowed parameter regions in our model that would address the ${}^7\text{Li}$ problem.

Compared to previous results, we note that leaving the reheating temperature and nonthermal abundance as free parameters opens up parameter space for higher ALP masses, which are otherwise excluded by cosmological constraints (see Fig. 1). Nonetheless, bounds on the decays of ALPs produced in SN1987A still provide relevant constraints in that region. This emphasises the importance of an analysis framework that allows for the inclusion of all relevant parameters and complementary constraints from astrophysics – both of which can be achieved with global-fitting tools such as GAMBIT.

Our findings also highlight the fact that measurements of the spectrum of the CMB from over 20 years ago can provide some of the leading cosmological constraints for ALPs with masses around the MeV scale. In particular, these constraints are stronger than those from the photodisintegration of light elements after BBN, which only constrain ALP masses $\gtrsim 5$ MeV. More accurate measurements from a future satellite mission such as PIXIE therefore promise substantially stronger constraints and the potential to discover weakly coupled MeV-scale ALPs via spectral distortions.

To conclude, we emphasise that our analysis was focused on a fairly restrictive set of ALP models, in which the ALP couplings to electrons and hadrons are assumed to be negligible. While this is a common assumption in the literature, there has been substantial, recent progress in understanding the effects of ALP-matter couplings on cosmology [118–121]. This research programme goes hand in hand with a detailed exploration of the importance of ALP-matter coupling for supernovae such as SN1987A [122–124]. Moreover, ALPs with

hadronic couplings have been proposed as a possible solution of the ${}^7\text{Li}$ problem [125]. It will therefore be very interesting to extend our setup and explore the interplay between such ALP models and observations, and the resulting constraints on more general ALP parameter spaces.

Acknowledgments

We thank Deanna Hooper, Lloyd Knox, Julien Lesgourgues, Kai Schmidt-Hoberg, and all members of the GAMBIT community for discussions. We are thankful to Sebastian Wild for collaboration in the early stages of this work. We are furthermore very grateful to Marco Hufnagel for help with ACROPOLIS, to Matteo Lucca and Nils Schöneberg for help with the spectral distortions code, and to Joerg Jaeckel and Pedro C. Malta for helpful discussions on Ref. [91] and letting us inspect the associated software code. CB received financial support from the Australian Research Council Discovery Projects DP180102209 and DP210101636. TEG, FK and PSt were funded by the Deutsche Forschungsgemeinschaft (DFG) through the Emmy Noether Grant No. KA 4662/1-1. SH was supported by the Alexander von Humboldt Foundation and the German Federal Ministry of Education and Research. DJEM is supported by an Ernest Rutherford Fellowship from the Science and Technologies Facilities Council (UK). PSc acknowledges funding support from the Australian Research Council under Future Fellowship FT190100814. PSt and SH thank FK for hospitality at RWTH Aachen, where part of this work was done. We thank PRACE for access to the Joliot-Curie HPC cluster at CEA. This work was performed using the Cambridge Service for Data Driven Discovery (CSD3), part of which is operated by the University of Cambridge Research Computing on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). The DiRAC component of CSD3 was funded by BEIS capital funding via STFC capital grants ST/P002307/1 and ST/R002452/1 and STFC operations grant ST/R00689X/1. DiRAC is part of the National e-Infrastructure. This work used the Scientific Compute Cluster at GWDG, the joint data centre of Max Planck Society for the Advancement of Science (MPG) and University of Göttingen. This article made use of `pippi` v2.1 [126].

A GAMBIT implementation

The computations and simulations performed in this work were done using the global fitting framework GAMBIT. In this appendix we document the modifications and additions we implemented for the purpose of this study. For full details on the structure of GAMBIT, we refer the reader to the GAMBIT core paper [32]. Further details on the CosmoBit extension can be found in Refs. [31, 127].

A.1 New models

The new models added to GAMBIT for this study are a generic parametrisation of a cosmological ALP, **GeneralCosmoALP**, two simplified ALP models **CosmoALP** and **CosmoALP_gg_tau**, a re-parametrisation of ΛCDM , **LCDM_zreio**, two new models for non-standard radiation, **etaBBN_rBBN_dNurBBN** and **rBBN_dNurBBN**, and two models of nuisance parameters, **cosmo_nuisance_BK15** and **cosmo_nuisance_SpectralDistortions**.

Various models of ALPs were already implemented in GAMBIT in Ref. [103], including the GeneralALP model, which sits at the top of the model hierarchy. We generalised this base model even more to take into account cosmological effects, such as the thermal abundance of ALPs and the reheating temperature of the Universe, introducing a new “parent model”

of GeneralALP named **GeneralCosmoALP**. We also added two other “child models” of **GeneralCosmoALP**: **CosmoALP** and **CosmoALP_gg_tau**. As ALPs may decay anytime before neutrino decoupling and the present day, translation functions are provided to the decaying DM model `DecayingDM_photon`, assuming that the primary decay mode is via photons. A “friendship relation” between these models ensures that energy injection to the CMB is taken into account in models with short ALP lifetimes.

GeneralCosmoALP: `gagg, gaee, gaN, fa, ma0, Tchi, beta, thetai, f0_thermal, T_R`

Generalization of the GeneralALP model from Ref. [103], with the addition of the thermal ALP abundance ξ (the `f0_thermal` parameter) and reheating temperature T_R .

CosmoALP: `Cagg, fa, ma0, thetai, f0_thermal, T_R`

Child model of **GeneralCosmoALP**, where the ALP-electron coupling and ALP-nucleon coupling are neglected, $g_{ae} = g_{aN} = 0$, and with temperature-independent mass, i.e. $\beta = 0$ and T_χ a large and irrelevant value. The ALP-photon coupling $g_{a\gamma}$ is replaced by its dimensionless equivalent $C_{a\gamma}$.

CosmoALP_gg_tau: `tau, fa, ma0, Tchi, beta, thetai, f0_thermal, T_R`

Child model of **GeneralCosmoALP** where the ALP-electron coupling and ALP-nucleon coupling are neglected, $g_{ae} = g_{aN} = 0$, and the ALP-photon coupling is replaced by the ALP lifetime τ_a .

The standard parametrisation of Λ CDM in GAMBIT includes the the baryon density ω_b , the DM density ω_{DM} , the Hubble constant H_0 , the CMB temperature T_{CMB} , and the optical depth at reionisation τ_{reio} . However, some of the external libraries we use, e.g. CLASS, internally use the redshift at reionisation z_{reio} instead of τ_{reio} . Typically obtaining z_{reio} from τ_{reio} is not an issue, but in the presence of energy injection at the time of reionisation this is not trivial. For this reason we introduce a new model, **LCDM_zreio**.

LCDM_zreio: `omega_b, omega_cdm, H0, T_cmb, z_reio`

Re-parametrisation of the Λ CDM model with the redshift at reionisation z_{reio} instead of the optical depth at reionisation τ_{reio} .

In GAMBIT, all non-standard radiation content that modifies the Hubble rate is parametrised by the five parameter model `etaBBN_rBBN_rCMB_dNurBBN_dNurCMB` and children thereof, cf. Ref. [31, appendix B.4]. For this analysis we have implemented two additional models, **etaBBN_rBBN_dNurBBN** and **rBBN_dNurBBN**, which allow for scenarios in which the modifications to N_{eff} at BBN are free model parameters that fully determine the respective values at recombination.

etaBBN_rBBN_dNurBBN: `dNur_BBN, r_BBN, eta_BBN`.

Parametrisation of non-standard radiation content with η_{BBN} , the baryon-to-photon ratio at the end of BBN; $\Delta N_{\text{ur, BBN}}$, the ultra-relativistic degrees of freedom at the end of BBN; and $r_{\nu, \text{BBN}}$, the neutrino temperature ratio at the end of BBN. A translation function is provided to map these parameters to those of `etaBBN_rBBN_rCMB_dNurBBN_dNurCMB`, requiring the new capability `Neff_evolution`, defined below.

rBBN_dNurBBN: `dNur_BBN`, `r_BBN`

Child model of `etaBBN_rBBN_dNurBBN`, where η_{BBN} is inferred from the baryon-to-photon ratio today η_0 , and thus the translation depends on the new capability `eta_ratio`, defined below.

Lastly, we introduced nuisance parameter models to include the systematic uncertainties of various experiments whose likelihoods are made available by MontePython. The nuisance model for the BICEP/Keck Array was updated to include the likelihood for the BK15 results [128] in the model `cosmo_nuisance_BK15`. Additionally, as this work is the first introduction of CMB spectral distortions in GAMBIT (cf. Section 3.3), we implemented a model to include the nuisance parameters for the COBE/FIRAS and PIXIE experiments: `cosmo_nuisance_SpectralDistortions`.

cosmo_nuisance_BK15: `BBbetadust`, `BBbetasync`, `BBdust`, `BSsync`, `BBalphadust`, `BBTdust`, `BBalphasync`, `BBdustsynccorr`, `EEtoBB_dust`, `EEtoBB_sync`, `Delta_dust`, `Delta_sync`, `gamma_corr`, `gamma_95`, `gamma_150`, `gamma_220`

Extension of the existing `cosmo_nuisance_BK14` model including additional systematics.

cosmo_nuisance_SpectralDistortions: `sd_delta_T`, `sd_T_D`, `sd_beta_D`, `sd_A_D`, `sd_T_CIB`, `sd_beta_CIB`, `sd_A_CIB`, `sd_alpha_sync`, `sd_omega_sync`, `sd_A_sync`, `sd_T_e`, `sd_EM`, `sd_nu_p_spin`, `sd_A_spin`, `sd_A_C0`, `sd_y_reio_nuisance`

The spectral distortion nuisance model contains parameters that take into account various uncertainties on the COBE/FIRAS and PIXIE experiments.

A.2 New capabilities provided in CosmoBit and DarkBit

In addition to the new models introduced for this study, we also added new capabilities, and modified existing ones, to work with the ALP models. Most of the new module functions that provide these capabilities were added to CosmoBit, as they had to do with the cosmological implications of ALPs, but one was added to DarkBit, as it deals with late-time astrophysical constraints. The full list of new capabilities, along with the module functions that provide them and their dependencies and backend requirements, can be seen in Table 3. We give a quick summary of the most important ones below.

A few of the newly added capabilities involve computing properties of the ALPs that can be derived from the model parameters. The capability `lifetime` computes the ALP lifetime as a function of the mass and coupling. The `minimum_abundance` capability computes the minimal freeze-in abundance of ALPs, produced via Primakoff processes, either analytically or numerically (by reading in tables pre-computed with micrOMEGAs), and `minimum_fraction` computes the fraction of the total abundance of dark matter constituted by the minimal freeze-in abundance. The `total_DM_abundance` is the total abundance of ALPs, produced by realignment and freeze-in, and `DM_fraction` provides the fraction of the total DM abundance constituted by ALPs (assuming no ALP decays).

In decaying DM models such as ours, the baryon-to-photon ratio at BBN η_{BBN} , the neutrino-to-photon temperature ratio r_ν , and the number of ultra-relativistic degrees of freedom at recombination ΔN_{eff} , are provided by the capability `external_dNeff_etaBBN`. These values are returned as a `std::map`, and are can be accessed individually using the capabilities `eta_ratio` for η_{BBN} and `Neff_evolution` for both r_ν and ΔN_{eff} . Details of these computations can be found in Appendix B.

Table 3: Capabilities added to CosmoBit and DarkBit associated with this study. Here, `map_str_dbl` is a `typedef` for `std::map<std::string,double>`, `map_str_int` for `std::map<std::string,int>` and `[pd]` is shorthand for `photodisintegration`. All functions are contained in CosmoBit except for `calc_PhotonFluence_SN1987A_Decay`, which is part of DarkBit.

Capability	Function (type)	Dependencies [type] / Backend reqs [type (args)]
<code>lifetime</code>	<code>lifetime_ALP_agg (double)</code>	
<code>minimum_abundance</code>	<code>minimum_abundance_ALP_analytical (double)</code> <code>minimum_abundance_ALP_numerical (double)</code>	
<code>minimum_fraction</code>	<code>minimum_fraction_ALP (double)</code>	<code>minimum_abundance [double]</code>
<code>DM_fraction</code>	<code>DM_fraction_ALP (double)</code>	<code>minimum_fraction [double]</code> <code>lifetime [double]</code> <code>RD_oh2 [double]</code>
<code>total_DM_abundance</code>	<code>total_DM_abundance_ALP (double)</code>	<code>DM_fraction [double]</code>
<code>external_dNeff_etaBBN</code>	<code>compute_dNeff_etaBBN_ALP (map_str_dbl)</code>	<code>total_DM_abundance [double]</code> <code>lifetime [double]</code>
<code>eta_ratio</code>	<code>eta_ratio_ALP (double)</code>	<code>external_dNeff_etaBBN [double]</code>
<code>Neff_evolution</code>	<code>Neff_evolution_ALP (map_str_dbl)</code>	<code>external_dNeff_etaBBN [double]</code>
<code>tau_reio</code>	<code>get_tau_reio_classy (double)</code>	<code>class_get_tau_reio [double ()]</code>
<code>z_reio</code>	<code>get_z_reio_classy (double)</code>	<code>class_get_z_reio [double ()]</code>
<code>H_at_z</code>	<code>get_H_at_z_classy (daFunk::Funk)</code>	<code>class_get_Hz [double (double)]</code>
<code>time_at_z</code>	<code>get_time_at_z_classy (daFunk::Funk)</code>	<code>class_get_tz [double (double)]</code>
<code>age_universe</code>	<code>get_age_universe_from_time_at_z (double)</code>	<code>time_at_z [daFunk::Funk]</code>
<code>primordial_abundances_BBN</code>	<code>compute_primordial_abundances_BBN (BBN_container)</code>	<code>AterBBN_Input [map_str_dbl]</code> <code>call_nucl_err [int (map_str_dbl, double, double)]</code> <code>call_nucl [int (map_str_dbl,double)]</code> <code>get_NNUC [size_t ()]</code> <code>get_abund_map_AlterBBN [map_str_int ()]</code>
<code>primordial_abundances</code>	<code>primordial_abundances (BBN_container)</code> <code>primordial_abundances_decayingDM (BBN_container)</code>	<code>primordial_abundances_BBN [BBN_container]</code> <code>set_input_params [void (bool,int,int,double)]</code> <code>abundance_[pd]_decay [void (double*,...)]</code>
<code>PhotonFluence_SN1987A_Decay</code>	<code>calc_PhotonFluence_SN1987A_Decay (double)</code>	

As mentioned above in A.1, we have defined a re-parametrisation of the Λ CDM model in GAMBIT to more closely match the parametrisation in CLASS. Along with the new model, a few capabilities are defined to provide useful quantities. `tau_reio` and `z_reio` extract the values of the optical depth and redshift of reionisation from CLASS. The capabilities `H_at_z` and `time_at_z` return the Hubble parameter and time since the Big Bang from CLASS as `daFunk` functions of redshift (see [102] for more information on the `daFunk` library). Lastly `age_universe` computes the age of the Universe at a given redshift with CLASS.

To take into account the effects of the photodisintegration of light elements due to the EM

spectrum injected by the ALP decays, we have restructured the capabilities relating to BBN. The capability `primordial_abundances_BBN` replaces the old `BBN_abundances` capability, which computes the abundances after BBN using `AlterBBN`. The new capability `primordial_abundances` computes the abundances at the present time. Two module functions provide this capability, `primordial_abundances`, which simply returns the abundances unchanged from BBN, and `primordial_abundances_decayingDM`, which uses `ACROPOLIS` to modify the computed abundances at BBN in order to include the effects of photodisintegration.

Finally, the only new capability added to `DarkBit` is `PhotonFluence_SN1987A_Decay`, which computes the γ fluence from the supernova SN1987A as a result of the decay of ALPs into photons.

A.3 New and updated backend interfaces

`GAMBIT` provides interfaces to a large collection of external tools, known as backends. The backend interfaces are documented along with the modules that use them [31, 102, 129–132]. Backend interfaces can also be auto-generated [133]. For this analysis we have updated some of the backends reported in [31] to their latest version, such as `CLASS` 3.1 [41] and `MontePython` 3.5 [86]. The updated version of `CLASS` brings along the possibility of including spectral distortions, as described in Section 3.3, which also benefits from the new version of `MontePython`, as it comes with more likelihoods, including `COBE/FIRAS` and a mock `PIXIE`-like likelihood. Furthermore, we have modified the interface to `AlterBBN` [54, 55] in order to allow the computation of abundances without uncertainties, as well as to return three values of the abundances, corresponding to the central, high and low values computed by `AlterBBN`.

The only absolutely new backend interface implemented for this study is to the external tool `ACROPOLIS` 1.2.1 [50]. Written in Python, `ACROPOLIS` computes the effects of photodisintegration of light elements on pre-computed abundances. In addition to the light element abundances, `ACROPOLIS` takes as input the mass of the decaying particle, m_a , its lifetime τ_a , the ratio of the number density of the decaying particle today to that of photons, $N_{0,a}$, and the branching fractions to electrons and photons, respectively. From this set of input parameters, the backend convenience function `abundance_photodisintegration_decay` performs various transformations on the input and output abundances and covariance matrix, and returns the final abundances after photodisintegration, and their associated covariance matrix.

The light element abundances in `ACROPOLIS` are taken as absolute abundances, as opposed to the treatment in `GAMBIT` and `AlterBBN`, where they are treated as ratios over the hydrogen abundance, $r_i = Y_i/Y_H$. Hence before feeding the abundances to `ACROPOLIS`, they are modified accordingly. Furthermore, the abundances computed by `AlterBBN` include the late-time decays of ${}^3\text{H}$ and ${}^7\text{Be}$. As these happen after photodisintegration, and are in fact automatically added by `ACROPOLIS`, we revert the ${}^3\text{H}$ and ${}^7\text{Be}$ decays when converting to absolute abundances for sending to `ACROPOLIS`. Conversely, before returning the final abundances, they are converted back to ratios over Y_H .

`ACROPOLIS` takes the abundances as a fixed-size array ordered as $\{n, p, D, {}^3\text{H}, {}^3\text{He}, {}^4\text{He}, {}^6\text{Li}, {}^7\text{Li}, {}^7\text{Be}\}$ and, along with the input parameters, computes a transfer matrix \mathbf{M} , which provides the abundances after photodisintegration as $\mathbf{Y}_f = \mathbf{M}\mathbf{Y}_i$. Besides the computation of the final abundances, this transfer matrix is used to propagate the covariance matrix as

$$\Sigma_f = (\mathbf{J}_2\mathbf{M}\mathbf{J}_1)\Sigma_i(\mathbf{J}_2\mathbf{M}\mathbf{J}_1)^T, \quad (\text{A.1})$$

where \mathbf{J}_1 is the Jacobian of the transformation from ratios over Y_H to absolute abundances, and \mathbf{J}_2 the converse. These Jacobians are different because the former includes the subtraction of late-time decays, whereas these are not included in the latter, but rather in the transfer matrix \mathbf{M} .

Lastly, there are few internal options in ACROPOLIS that determine the precision and speed of the computations, among others. These can be provided as options in the YAML file, which can be seen below, including three configurations for default, fast and aggressive settings. The backend convenience function `set_input_params` ensures that the given options are passed to the relevant ACROPOLIS modules.

```
# Options for ACROPOLIS
- capability: primordial_abundances
  function: primordial_abundances_decayingDM
  options:
    verbose: false
    NE_pd: 30 # default: 150, fast: 75, aggressive: 30
    NT_pd: 10 # default: 50, fast: 25, aggressive: 10
    eps: 1e-1 # default: 1e-3, fast: 1e-2, aggressive: 1e-1
```

B Details of implementation of entropy injection between BBN and recombination

The evolution of the photon temperature including photon production through ALP decays $a \rightarrow \gamma\gamma$ is given by Eq. (3.4). It is important to note that the functions $n_a(t)$ and $H(t)$ in Eq. (3.4) both have an implicit dependence on the photon temperature T_γ and neutrino temperature T_ν . Both temperatures affect the energy density in radiation ρ_{rad} , see Eq. (3.1), which in turn affects the expansion rate $H(t)$ and the scale factor $R(t)$ entering $n_a(t)$. Self-consistently integrating Eq. (3.4) therefore requires taking the implicit temperature dependence into account. Doing so can, however, lead to numerical instabilities. In order to circumvent these complications, we take an alternative approach: instead of solving Eq. (3.4) only once, while respecting the implicit dependence of $H(t)$ and $n_a(t)$ on T_γ , we solve this equation several times while neglecting the changes in $H(t)$ and $R(t)$ due to the changes in T_γ . After each iteration $H(t)$ and $R(t)$ are updated and the differential equation is solved again. This means that in the i th iteration, the differential equation is given by

$$\frac{dT_\gamma^{(i)}}{dt} = \frac{15}{4\pi^2} \frac{m_a n_a^{(i-1)}(t)}{\tau_a} \frac{1}{(T_\gamma^{(i)})^3} - H^{(i-1)}(t) T_\gamma^{(i)}, \quad (\text{B.1})$$

where $H^{(i-1)}(t)$ and $n_a^{(i-1)}(t)$ are derived assuming the photon temperature of the previous iteration $T_\gamma^{(i-1)}$. This procedure is repeated until T_γ has converged, i.e. $T_\gamma^{(i)}(t) = T_\gamma^{(i-1)}(t)$.

B.1 Numerical implementation

Here we describe the numerical implementation of this procedure. We create a look-up table of the photon temperature $T_\gamma(t)$, the neutrino temperature $T_\nu(t)$, the Hubble rate $H(t)$, as well as the natural logarithm of the scale factor $\ln R(t)$, which enter $n_a(t)$ as given by Eq. (3.5).¹⁴

¹⁴In the scenario studied in this paper, the neutrino temperature can be derived through Eq. (3.6) such that there would be no need to include $T_\nu(t)$ in this table, as it can be directly derived through $\ln R(t)$. We include this, however, to allow for easy generalisation to scenarios in which the neutrino temperature is also modified.

When this table is initialised ($i = 0$), it has logarithmic spacing in time between t_{ini} and t_{fin} . Furthermore, we assume that there are no modifications to the photon temperature and neutrino temperature, meaning that initially the effective number of relativistic degrees of freedom $g_*(T)$ is assumed to be constant in the temperature range of interest. In this case $H = \dot{R}/R \sim T_\gamma^2$ and $T_\gamma \sim 1/R$ such that the initial Hubble rate $H^{(0)}(t)$ is given by

$$H^{(0)}(t) = \frac{1}{2t}. \quad (\text{B.2})$$

Using the right-hand side of Eq. (2.4), the initial photon temperature $T_\gamma^{(0)}(t)$ is calculated from $H^{(0)}(t)$ and the initial neutrino temperature is given by assuming instant decoupling $T_\nu^{(0)} = (4/11)^{1/3} T_\gamma^{(0)}$. The last column of the table $\ln a^{(0)}(t)$ is initialised via

$$\ln R^{(0)}(t) = \int_{t_{\text{ini}}}^t dt' H^{(0)}(t'). \quad (\text{B.3})$$

After Eq. (B.1) is solved the table has to be updated in order to be self-consistent. In particular, as $T_\gamma(t)$ is modified, $H(t)$ is modified as well. Moreover, it is important to note that the constant column in this table is not the time t but the scale factor or its natural logarithm $\ln R(t)$. Hence, once the column containing H is updated by using the respective values of T_γ (and T_ν) after solving the differential equation, the column containing the time coordinate t has to be updated as well by using $H = d \ln R/dt$, such that t becomes a function of $\ln R$ and is given by

$$t(\ln R) = \int_{\ln R_0}^{\ln R} d \ln R' \frac{1}{H(\ln R')}. \quad (\text{B.4})$$

Once H and t have been self-consistently updated, the differential equation can be solved again.

C Details of the SN1987A limits from ALP decay

Here we give the full details of our SN1987A likelihood. While a description of the procedure and some of the relevant equations was already provided in the literature [91, 134, 135], we extend the discussion by adding more details on the reference frame transformations and Monte Carlo simulations.

C.1 Reference frame transformations

In Fig. 4, we show the ALP a emerging from SN1987A and decaying into two photons γ_1 and γ_2 with angles α_1 and α_2 . For convenience, let the x -axis be parallel to the direction of the emergent ALP. We denote the ALP rest frame by “0” and the lab frame, i.e. the rest frame of the SMM satellite, by “L.” The Lorentz boost into the lab frame, for dimensionless velocity β and $\gamma \equiv 1/\sqrt{1-\beta^2}$, is then given by

$$\Lambda = \begin{pmatrix} \gamma & \beta\gamma & 0 & 0 \\ \beta\gamma & \gamma & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \quad (\text{C.1})$$

Boosting the 4-momentum of an ALP from the rest frame to the lab frame gives

$$p_{a,0}^\mu = \begin{pmatrix} m_a \\ 0 \\ 0 \\ 0 \end{pmatrix} \mapsto p_{a,L}^\mu = \Lambda^{\mu\nu} (p_{a,0})_\nu = \begin{pmatrix} \gamma m_a \\ \beta\gamma m_a \\ 0 \\ 0 \end{pmatrix}. \quad (\text{C.2})$$

The two decay photons are emitted back-to-back in the ALP rest frame, which can be used to boost them into the lab frame:

$$p_{1/2,0}^\mu = \frac{m_a}{2} \begin{pmatrix} 1 \\ \pm \cos \phi \sin \theta \\ \pm \sin \phi \sin \theta \\ \pm \cos \theta \end{pmatrix} \mapsto p_{1/2,L}^\mu = \frac{\gamma m_a}{2} \begin{pmatrix} 1 \pm \beta \cos \phi \sin \theta \\ \beta \pm \cos \phi \sin \theta \\ \pm \frac{1}{\gamma} \sin \phi \sin \theta \\ \pm \frac{1}{\gamma} \cos \theta \end{pmatrix}. \quad (\text{C.3})$$

Consequently, the two photon energies E_1 and E_2 in the lab frame are

$$E_{1/2,L} = \frac{\gamma m_a}{2} (1 \pm \beta \cos \phi \sin \theta) = \frac{E_a}{2} (1 \pm \beta \cos \phi \sin \theta), \quad (\text{C.4})$$

while the photon angles α_1 and α_2 in the lab frame are defined geometrically as the angle between the 3-momenta and the x -direction \mathbf{e}_x such that

$$\cos \alpha_{1/2,L} = \frac{\mathbf{e}_x \cdot \mathbf{p}_{1/2}}{|\mathbf{e}_x| |\mathbf{p}_{1/2}|} = \frac{E_a}{2} \frac{\beta \pm \cos \phi \sin \theta}{|\mathbf{p}_{1/2}|}. \quad (\text{C.5})$$

Calculating the angles in Eq. (C.5) requires some algebra, in particular we need that

$$\frac{2}{E_a} |\mathbf{p}_{1/2}| = |1 \pm \beta \cos \phi \sin \theta| = 1 \pm \beta \cos \phi \sin \theta, \quad (\text{C.6})$$

where the absolute value in the last line can be dropped because the absolute values of β , $\cos \phi$, and $\sin \theta$ are all smaller than unity. The angles are therefore given by

$$\alpha_{1/2,L} = \arccos \left(\frac{\beta \pm \cos \phi \sin \theta}{1 \pm \beta \cos \phi \sin \theta} \right). \quad (\text{C.7})$$

Equation (C.7) reveals additional rotational symmetry in the problem because only the product $\cos \phi \sin \theta$ is important; one could thus e.g. set $\theta = \pi$ for simplicity.

C.2 Monte Carlo procedure

For each $(m_a, g_{a\gamma})$ pair we simulate 10^7 events. This numbers reasonable since the naïve photon fluence in the parameter region of interest is of order 10^6 cm^{-2} while the 3σ fluence limit is about 2 cm^{-2} . For example, ALPs with $m_a = 1 \text{ MeV}$ and $g_{a\gamma} = 10^{-11} \text{ GeV}^{-1}$ are just outside of the excluded region; the corresponding acceptance fraction of the MC procedure is about 2×10^{-5} , or 203 ± 14 events out of 10^7 .

We first draw ALP energies from the spectrum in Eq. (3.17) by numerically calculating the inverse of the (normalised) spectrum's cumulative distribution. Similarly, we draw the decay lengths L_a from an exponential distribution, with mean value parameter given by Eq. (3.18). For ALPs to yield detectable decays, they have to:

- a) decay at least 10^{10} m away from SN1987A (the orange shell around SN1987A in Fig. 4), as photons from decays closer to the SN will be absorbed or deflected by its turbulent environment, and
- b) decay before reaching the Earth, as almost no photons will reach the detector beyond this point (cf. discussion in Ref. [91]).

For photons that satisfy these conditions, we draw a random emission angle for one of the two photons in the ALP rest frame from uniform distributions: $\phi \sim \mathcal{U}(0, 2\pi)$ and $\cos\theta \sim \mathcal{U}(-1, 1)$. The angles ϕ and θ of the other photon are the negative of the angles of the first photon, as the two photons are emitted back-to-back. Then, for each of the two photons, we calculate the energies $E_{1/2,L}$ according to Eq. (C.6) in the lab frame and check that these energies are within the detector energy window of 25–100 MeV.

Finally, for the photons also making this cut, we calculate the photon angles $\alpha_{1/2}$ in the lab frame according to Eq. (C.7). This allows us to compute the total length of their path, $L_{1/2}$, to the detector and hence their arrival times $t_{1/2}$, which are [91, Sec. II.C]

$$L_{\gamma,1/2} \equiv -L_a \cos(\alpha_{1/2}) \sqrt{d_{\text{SN}}^2 - L_a^2 \sin^2(\alpha_{1/2})}, \quad (\text{C.8})$$

$$t_{1/2} = L_{1/2} = L_a/\beta + L_{\gamma,1/2} - d_{\text{SN}}. \quad (\text{C.9})$$

Demanding that photons arrive earlier than 223s after the arrival of the first neutrino is the last check in the Monte Carlo procedure.

C.3 Results

We tabulate the calculated fluence on a logarithmic grid of masses from 1 keV to 1 GeV and ALP-photon couplings from $10^{-13} \text{ GeV}^{-1}$ to 10^{-3} GeV^{-1} . We use a grid spacing of 0.05 dex in both dimensions, leading to a grid of 24,000 points. We chose this grid in order to close the gap between SN1987A and Horizontal Branch star limits [e.g. 91]. The resulting table is provided in GAMBIT as `DarkBit/data/SN1987A_DecayFluence.dat`.

The results are shown in Fig. 4, where we compare the excluded region from our implementation with that of Ref. [91]. Our contours are smoother as we perform full Monte Carlo simulations without the high-mass cut used in Ref. [91]. Apart from this difference, we get good qualitative agreement with the previous bounds. However, note that we typically find that our calculated flux is a factor $\lesssim 2$ larger than Ref. [91] – in line with what Ref. [124] observe using the non-probabilistic integration methods from Ref. [136]. After inspecting the code of Ref. [91] and its output, we suspect that apart from the wide binning used in Ref. [91] there might also be some “double counting” in the implementation of the factors considered in their Eq. (14).

C.4 Technical details

Our scripts, written in Python v3.7.5 and using the `numpy`, `scipy`, `time`, and `mpi4py` packages, are available at <https://github.com/marie-lecroq/ALP-fluence-calculation>. The main script `functions.py` contains the function `expected_photon_fluence`, which returns the photon fluence (in units of cm^{-2}) given an ALP mass `m` (in units of eV) and ALP-photon coupling `g` (in units of GeV^{-1}). We also included the script `run_analysis_mpi.py` to run our algorithm in parallel using MPI.

References

- [1] R. Diamanti, L. Lopez-Honorez, O. Mena, S. Palomares-Ruiz, and A. C. Vincent, *Constraining Dark Matter Late-Time Energy Injection: Decays and P-Wave Annihilations*, *JCAP* **02** (2014) 017, [[arXiv:1308.2578](#)].
- [2] T. R. Slatyer and C.-L. Wu, *General Constraints on Dark Matter Decay from the Cosmic Microwave Background*, *Phys. Rev. D* **95** (2017) 023010, [[arXiv:1610.06933](#)].
- [3] V. Poulin, J. Lesgourgues, and P. D. Serpico, *Cosmological constraints on exotic injection of electromagnetic energy*, *JCAP* **03** (2017) 043, [[arXiv:1610.10051](#)].
- [4] W. Hu and J. Silk, *Thermalization constraints and spectral distortions for massive unstable relic particles*, *Phys. Rev. Lett.* **70** (1993) 2661–2664.
- [5] M. Lucca, N. Schöneberg, D. C. Hooper, J. Lesgourgues, and J. Chluba, *The synergy between CMB spectral distortions and anisotropies*, *JCAP* **2020** (2020) 026, [[arXiv:1910.04619](#)].
- [6] V. Poulin and P. D. Serpico, *Nonuniversal BBN bounds on electromagnetically decaying particles*, *Phys. Rev. D* **91** (2015) 103007, [[arXiv:1503.04852](#)].
- [7] M. Hufnagel, K. Schmidt-Hoberg, and S. Wild, *BBN constraints on MeV-scale dark sectors. Part II. Electromagnetic decays*, *JCAP* **1811** (2018) 032, [[arXiv:1808.09324](#)].
- [8] P. F. Depta, M. Hufnagel, and K. Schmidt-Hoberg, *Updated BBN constraints on electromagnetic decays of MeV-scale particles*, *JCAP* **04** (2021) 011, [[arXiv:2011.06519](#)].
- [9] X.-L. Chen and M. Kamionkowski, *Particle decays during the cosmic dark ages*, *Phys. Rev. D* **70** (2004) 043502, [[astro-ph/0310473](#)].
- [10] J. E. Kim, *Light pseudoscalars, particle physics and cosmology.*, *Phys. Rep.* **150** (1987) 1–177.
- [11] J. Jaeckel and A. Ringwald, *The Low-Energy Frontier of Particle Physics*, *Annual Review of Nuclear and Particle Science* **60** (2010) 405–437, [[arXiv:1002.0329](#)].
- [12] M. Cicoli, M. Goodsell, and A. Ringwald, *The type IIB string axiverse and its low-energy phenomenology*, *JHEP* **10** (2012) 146, [[arXiv:1206.0819](#)].
- [13] S. Weinberg, *A new light boson?*, *Physical Review Letters* **40** (1978) 223–226.
- [14] F. Wilczek, *Problem of strong P and T invariance in the presence of instantons*, *Physical Review Letters* **40** (1978) 279–282.
- [15] I. Brivio, M. B. Gavela, *et. al.*, *ALPs Effective Field Theory and Collider Signatures*, *Eur. Phys. J. C* **77** (2017) 572, [[arXiv:1701.05379](#)].
- [16] M. Bauer, M. Neubert, S. Renner, M. Schnubel, and A. Thamm, *The Low-Energy Effective Theory of Axions and ALPs*, *JHEP* **04** (2021) 063, [[arXiv:2012.12272](#)].
- [17] E. Massó and R. Toldrà, *Light spinless particle coupled to photons*, *Phys. Rev. D* **52** (1995) 1755–1763, [[hep-ph/9503293](#)].
- [18] E. Massó and R. Toldrà, *New constraints on a light spinless particle coupled to photons*, *Phys. Rev. D* **55** (1997) 7967–7969, [[hep-ph/9702275](#)].
- [19] D. Cadamuro and J. Redondo, *Cosmological bounds on pseudo Nambu-Goldstone bosons*, *JCAP* **2012** (2012) 032, [[arXiv:1110.2895](#)].
- [20] M. Millea, L. Knox, and B. D. Fields, *New bounds for axions and axion-like particles with keV-GeV masses*, *Phys. Rev. D* **92** (2015) 023010, [[arXiv:1501.04097](#)].
- [21] P. Arias, D. Cadamuro, *et. al.*, *WISPy cold dark matter*, *JCAP* **6** (2012) 013, [[arXiv:1201.5902](#)].
- [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].
- [23] M. Giannotti, I. Irastorza, J. Redondo, and A. Ringwald, *Cool WISPs for stellar cooling excesses*, *JCAP* **5** (2016) 057, [arXiv:1512.08108].
- [24] M. Bauer, M. Neubert, and A. Thamm, *Collider Probes of Axion-Like Particles*, *JHEP* **12** (2017) 044, [arXiv:1708.00443].
- [25] M. J. Dolan, T. Ferber, C. Hearty, F. Kahlhoefer, and K. Schmidt-Hoberg, *Revised constraints and Belle II sensitivity for visible and invisible axion-like particles*, *JHEP* **12** (2017) 094, [arXiv:1709.00009]. [Erratum: *JHEP* **03**, 190 (2021)].
- [26] D. J. E. Marsh, *Axion cosmology*, *Phys. Rep.* **643** (2016) 1–79, [arXiv:1510.07633].
- [27] I. G. Irastorza and J. Redondo, *New experimental approaches in the search for axion-like particles*, *Progress in Particle and Nuclear Physics* **102** (2018) 89–159, [arXiv:1801.08127].
- [28] C. O’Hare, “AxionLimits: Data, plots and code for constraints on axions, axion-like particles, and dark photons.” Zenodo, 2020. Current version of the repository available on Github at <https://github.com/cajohare/AxionLimits>, [DOI:10.5281/zenodo.3932430].
- [29] P. F. Depta, M. Hufnagel, and K. Schmidt-Hoberg, *Robust cosmological constraints on axion-like particles*, *JCAP* **2020** (2020) 009, [arXiv:2002.08370].
- [30] M. Bolz, A. Brandenburg, and W. Buchmuller, *Thermal production of gravitinos*, *Nucl. Phys. B* **606** (2001) 518–544, [hep-ph/0012052]. [Erratum: *Nucl. Phys. B* **790**, 336 (2008)].
- [31] GAMBIT Cosmology Workgroup: J. J. Renk, P. Stöcker, *et. al.*, *CosmoBit: A GAMBIT module for computing cosmological observables and likelihoods*, *JCAP* **02** (2021) 022, [arXiv:2009.03286].
- [32] GAMBIT Collaboration: P. Athron, C. Balázs, *et. al.*, *GAMBIT: The Global and Modular Beyond-the-Standard-Model Inference Tool*, *Eur. Phys. J. C* **77** (2017) 784, [arXiv:1705.07908]. Addendum in [137].
- [33] A. Kvellestad, P. Scott, and M. White, *GAMBIT and its Application in the Search for Physics Beyond the Standard Model*, *Prog. Part. Nuc. Phys.* **113** (2020) 103769, [arXiv:1912.04079].
- [34] A. Ayala, I. Domínguez, M. Giannotti, A. Mirizzi, and O. Straniero, *Revisiting the Bound on Axion-Photon Coupling from Globular Clusters*, *Phys. Rev. Lett.* **113** (2014) 191302, [arXiv:1406.6053].
- [35] A. Boyarsky, A. Neronov, O. Ruchayskiy, M. Shaposhnikov, and I. Tkachev, *Strategy for Searching for a Dark Matter Sterile Neutrino*, *Phys. Rev. Lett.* **97** (2006) 261302, [astro-ph/0603660].
- [36] A. Boyarsky, O. Ruchayskiy, and M. Shaposhnikov, *The Role of Sterile Neutrinos in Cosmology and Astrophysics*, *Annual Review of Nuclear and Particle Science* **59** (2009) 191–214, [arXiv:0901.0011].
- [37] J. W. Foster, M. Kongsore, *et. al.*, *Deep Search for Decaying Dark Matter with XMM-Newton Blank-Sky Observations*, *Phys. Rev. Lett.* **127** (2021) 051101, [arXiv:2102.02207].
- [38] GAMBIT Collaboration, *Supplementary Data: Cosmological constraints on decaying axion-like particles: a global analysis.*, (2022), <https://zenodo.org/record/6573347>.
- [39] J. Jaeckel, J. Redondo, and A. Ringwald, *3.55 keV hint for decaying axionlike particle dark matter*, *Phys. Rev. D* **89** (2014) 103511, [arXiv:1402.7335].
- [40] P. F. de Salas, M. Lattanzi, *et. al.*, *Bounds on very low reheating scenarios after Planck*, *Phys. Rev. D* **92** (2015) 123534, [arXiv:1511.00672].
- [41] D. Blas, J. Lesgourgues, and T. Tram, *The Cosmic Linear Anisotropy Solving System (CLASS) II: Approximation schemes*, *JCAP* **1107** (2011) 034, [arXiv:1104.2933].

- [42] Planck: N. Aghanim *et. al.*, *Planck 2018 results. VI. Cosmological parameters*, *Astron. Astrophys.* **641** (2020) A6, [[arXiv:1807.06209](#)].
- [43] Planck: N. Aghanim *et. al.*, *Planck 2018 results. V. CMB power spectra and likelihoods*, *Astron. Astrophys.* **641** (2020) A5, [[arXiv:1907.12875](#)].
- [44] P. Stöcker, M. Krämer, J. Lesgourgues, and V. Poulin, *Exotic energy injection with ExoCLASS: Application to the Higgs portal model and evaporating black holes*, *JCAP* **1803** (2018) 018, [[arXiv:1801.01871](#)].
- [45] GAMBIT: P. Athron *et. al.*, *Thermal WIMPs and the scale of new physics: global fits of Dirac dark matter effective field theories*, *Eur. Phys. J. C* **81** (2021) 992, [[arXiv:2106.02056](#)].
- [46] J. J. Bennett, G. Buldgen, M. Drewes, and Y. Y. Y. Wong, *Towards a precision calculation of the effective number of neutrinos N_{eff} in the Standard Model: the QED equation of state*, *JCAP* **2020** (2020) 003, [[arXiv:1911.04504](#)].
- [47] K. Akita and M. Yamaguchi, *A precision calculation of relic neutrino decoupling*, *JCAP* **2020** (2020) 012, [[arXiv:2005.07047](#)].
- [48] J. Froustey, C. Pitrou, and M. C. Volpe, *Neutrino decoupling including flavour oscillations and primordial nucleosynthesis*, *JCAP* **2020** (2020) 015, [[arXiv:2008.01074](#)].
- [49] J. J. Bennett, G. Buldgen, *et. al.*, *Towards a precision calculation of N_{eff} in the Standard Model II: Neutrino decoupling in the presence of flavour oscillations and finite-temperature QED*, *JCAP* **2021** (2021) 073, [[arXiv:2012.02726](#)].
- [50] P. F. Depta, M. Hufnagel, and K. Schmidt-Hoberg, *ACROPOLIS: A generiC fRamework fOR Photodisintegration Of LIght elementS*, *JCAP* **03** (2021) 061, [[arXiv:2011.06518](#)].
- [51] Particle Data Group: P. A. Zyla *et. al.*, *Review of Particle Physics*, *Prog. Theor. Exp. Phys.* **083** (2020) C01.
- [52] J. Geiss and G. Gloeckler, *Isotopic Composition of H, HE and NE in the Protosolar Cloud*, *Space Sci. Rev.* **106** (2003) 3–18.
- [53] P. F. Depta, M. Hufnagel, K. Schmidt-Hoberg, and S. Wild, *BBN constraints on the annihilation of MeV-scale dark matter*, *JCAP* **1904** (2019) 029, [[arXiv:1901.06944](#)].
- [54] A. Arbey, *AlterBBN: A program for calculating the BBN abundances of the elements in alternative cosmologies*, *Comput. Phys. Commun.* **183** (2012) 1822–1831, [[arXiv:1106.1363](#)].
- [55] A. Arbey, J. Auffinger, K. P. Hickerson, and E. S. Jenssen, *AlterBBN v2: A public code for calculating Big-Bang nucleosynthesis constraints in alternative cosmologies*, *Comput. Phys. Commun.* **248** (2020) 106982, [[arXiv:1806.11095](#)].
- [56] B. D. Fields, *The primordial lithium problem*, *Ann. Rev. Nucl. Part. Sci.* **61** (2011) 47–68, [[arXiv:1203.3551](#)].
- [57] D. J. Fixsen, E. S. Cheng, *et. al.*, *The Cosmic Microwave Background Spectrum from the Full COBE FIRAS Data Set*, *ApJ* **473** (1996) 576, [[astro-ph/9605054](#)].
- [58] D. J. Fixsen, *The Temperature of the Cosmic Microwave Background*, *ApJ* **707** (2009) 916–920, [[arXiv:0911.1955](#)].
- [59] Y. B. Zeldovich and R. A. Sunyaev, *The Interaction of Matter and Radiation in a Hot-Model Universe*, *Ap&SS* **4** (1969) 301–316.
- [60] R. A. Sunyaev and Y. B. Zeldovich, *The interaction of matter and radiation in the hot model of the Universe, II*, *Ap&SS* **7** (1970) 20–30.
- [61] Y. B. Zeldovich, A. F. Illarionov, and R. A. Syunyaev, *The Effect of Energy Release on the Emission Spectrum in a Hot Universe*, *ZhETF* **68** (1972) 1217.
- [62] A. F. Illarionov and R. A. Siuniaev, *Comptonization, characteristic radiation spectra, and*

- thermal balance of low-density plasma*, *Soviet Ast.* **18** (1975) 413–419.
- [63] L. Danese and G. de Zotti, *Double Compton process and the spectrum of the microwave background*, *A&A* **107** (1982) 39–42.
- [64] C. Burigana, L. Danese, and G. de Zotti, *Formation and evolution of early distortions of the microwave background spectrum - A numerical study*, *A&A* **246** (1991) 49–58.
- [65] J. Chluba and R. A. Sunyaev, *The evolution of CMB spectral distortions in the early Universe*, *MNRAS* **419** (2012) 1294–1314, [[arXiv:1109.6552](#)].
- [66] J. Chluba, A. Kogut, *et. al.*, *Spectral Distortions of the CMB as a Probe of Inflation, Recombination, Structure Formation and Particle Physics*, *BAAS* **51** (2019) 184, [[arXiv:1903.04218](#)].
- [67] T. Yanagida and M. Yoshimura, *Resonant axion-photon conversion in the early universe*, *Physics Letters B* **202** (1988) 301–306.
- [68] A. Mirizzi, J. Redondo, and G. Sigl, *Constraining resonant photon-axion conversions in the early universe*, *JCAP* **2009** (2009) 001, [[arXiv:0905.4865](#)].
- [69] T. Higaki, K. Nakayama, and F. Takahashi, *Cosmological constraints on axionic dark radiation from axion-photon conversion in the early Universe*, *JCAP* **2013** (2013) 030, [[arXiv:1306.6518](#)].
- [70] H. Tashiro, J. Silk, and D. J. E. Marsh, *Constraints on primordial magnetic fields from CMB distortions in the axiverse*, *Phys. Rev. D* **88** (2013) 125024, [[arXiv:1308.0314](#)].
- [71] D. Ejlili and A. D. Dolgov, *CMB constraints on mass and coupling constant of light pseudoscalar particles*, *Phys. Rev. D* **90** (2014) 063514, [[arXiv:1312.3558](#)].
- [72] S. Mukherjee, R. Khatri, and B. D. Wandelt, *Polarized anisotropic spectral distortions of the CMB: galactic and extragalactic constraints on photon-axion conversion*, *JCAP* **2018** (2018) 045, [[arXiv:1801.09701](#)].
- [73] S. Mukherjee, R. Khatri, and B. D. Wandelt, *Constraints on non-resonant photon-axion conversion from the Planck satellite data*, *JCAP* **2019** (2019) 031, [[arXiv:1811.11177](#)].
- [74] S. Mukherjee, D. N. Spergel, R. Khatri, and B. D. Wandelt, *A new probe of axion-like particles: CMB polarization distortions due to cluster magnetic fields*, *JCAP* **2020** (2020) 032, [[arXiv:1908.07534](#)].
- [75] D. Cadamuro, S. Hannestad, G. Raffelt, and J. Redondo, *Cosmological bounds on sub-MeV mass axions*, *JCAP* **2011** (2011) 003, [[arXiv:1011.3694](#)].
- [76] A. S. Kompaneets, *The Establishment of Thermal Equilibrium between Quanta and Electrons*, *Soviet Journal of Experimental and Theoretical Physics* **4** (1957) 730–737.
- [77] A. P. Lightman, *Double Compton emission in radiation dominated thermal plasmas*, *ApJ* **244** (1981) 392–405.
- [78] J. Chluba, M. H. Abitbol, *et. al.*, *New horizons in cosmology with spectral distortions of the cosmic microwave background*, *Experimental Astronomy* **51** (2021) 1515–1554, [[arXiv:1909.01593](#)].
- [79] J. Chluba, *Refined approximations for the distortion visibility function and μ -type spectral distortions*, *MNRAS* **440** (2014) 2544–2563, [[arXiv:1312.6030](#)].
- [80] J. C. Mather, E. S. Cheng, *et. al.*, *Measurement of the Cosmic Microwave Background Spectrum by the COBE FIRAS Instrument*, *ApJ* **420** (1994) 439.
- [81] D. J. Fixsen, E. Dwek, J. C. Mather, C. L. Bennett, and R. A. Shafer, *The Spectrum of the Extragalactic Far-Infrared Background from the COBE FIRAS Observations*, *ApJ* **508** (1998) 123–128, [[astro-ph/9803021](#)].

- [82] J. Chluba, *Green's function of the cosmological thermalization problem*, *MNRAS* **434** (2013) 352–357, [[arXiv:1304.6120](#)].
- [83] J. Chluba, D. Nagai, S. Sazonov, and K. Nelson, *A fast and accurate method for computing the Sunyaev-Zel'dovich signal of hot galaxy clusters*, *MNRAS* **426** (2012) 510–530, [[arXiv:1205.5778](#)].
- [84] J. Chluba, E. Switzer, K. Nelson, and D. Nagai, *Sunyaev-Zeldovich signal processing and temperature-velocity moment method for individual clusters*, *MNRAS* **430** (2013) 3054–3069, [[arXiv:1211.3206](#)].
- [85] B. Audren, J. Lesgourgues, K. Benabed, and S. Prunet, *Conservative Constraints on Early Cosmology: an illustration of the Monte Python cosmological parameter inference code*, *JCAP* **1302** (2013) 001, [[arXiv:1210.7183](#)].
- [86] T. Brinckmann and J. Lesgourgues, *MontePython 3: boosted MCMC sampler and other features*, *Phys. Dark Univ.* **24** (2019) 100260, [[arXiv:1804.07261](#)].
- [87] A. Kogut, D. J. Fixsen, *et. al.*, *The Primordial Inflation Explorer (PIXIE): a nulling polarimeter for cosmic microwave background observations*, *JCAP* **2011** (2011) 025, [[arXiv:1105.2044](#)].
- [88] J. Chluba and D. Jeong, *Teasing bits of information out of the CMB energy spectrum*, *MNRAS* **438** (2014) 2065–2082, [[arXiv:1306.5751](#)].
- [89] N. Schöneberg, M. Lucca, and D. C. Hooper, *Constraining the inflationary potential with spectral distortions*, *JCAP* **2021** (2021) 036, [[arXiv:2010.07814](#)].
- [90] M. H. Abitbol, J. Chluba, J. C. Hill, and B. R. Johnson, *Prospects for Measuring Cosmic Microwave Background Spectral Distortions in the Presence of Foregrounds*, *Mon. Not. Roy. Astron. Soc.* **471** (2017) 1126–1140, [[arXiv:1705.01534](#)].
- [91] J. Jaeckel, P. C. Malta, and J. Redondo, *Decay photons from the axionlike particles burst of type II supernovae*, *Phys. Rev. D* **98** (2018) 055032, [[arXiv:1702.02964](#)].
- [92] E. L. Chupp, W. T. Vestrand, and C. Reppin, *Experimental limits on the radiative decay of SN 1987A neutrinos*, *Phys. Rev. Lett.* **62** (1989) 505–508.
- [93] D. J. Forrest, E. L. Chupp, *et. al.*, *The gamma ray spectrometer for the Solar Maximum Mission.*, *Sol. Phys.* **65** (1980) 15–23.
- [94] E. L. Chupp, W. T. Vestrand, and C. Reppin, *Experimental limits on the radiative decay of SN 1987A neutrinos*, *Phys. Rev. Lett.* **62** (1989) 505–508.
- [95] A. Payez, C. Evoli, *et. al.*, *Revisiting the SN1987A gamma-ray limit on ultralight axion-like particles*, *JCAP* **1502** (2015) 006, [[arXiv:1410.3747](#)].
- [96] K. Sato and H. Sato, *Higgs meson emission from a star and a constraint on its mass*, *Progress of Theoretical Physics* **54** (1975) 1564.
- [97] G. G. Raffelt, *Astrophysical methods to constrain axions and other novel particle phenomena*, *Phys. Rep.* **198** (1990) 1–113.
- [98] G. G. Raffelt, *Stars as Laboratories for Fundamental Physics: The Astrophysics of Neutrinos, Axions, and Other Weakly Interacting Particles*. University Of Chicago Press, Chicago & London, 1996.
- [99] A. Buzzoni, F. F. Pecci, R. Buonanno, and C. E. Corsi, *Helium abundance in globular clusters - The R-method*, *A&A* **128** (1983) 94–101.
- [100] G. G. Raffelt, *Core mass at the helium flash from observations and a new bound on neutrino electromagnetic properties*, *ApJ* **365** (1990) 559–568.
- [101] E. Aver, K. A. Olive, and E. D. Skillman, *The effects of He I $\lambda 10830$ on helium abundance*

- determinations, *JCAP* **7** (2015) 011, [[arXiv:1503.08146](#)].
- [102] GAMBIT Dark Matter Workgroup: T. Bringmann, J. Conrad, *et. al.*, *DarkBit: A GAMBIT module for computing dark matter observables and likelihoods*, *Eur. Phys. J. C* **77** (2017) 831, [[arXiv:1705.07920](#)].
- [103] S. Hoof, F. Kahlhoefer, P. Scott, C. Weniger, and M. White, *Axion global fits with Peccei-Quinn symmetry breaking before inflation using GAMBIT*, *JHEP* **03** (2019) 191, [[arXiv:1810.07192](#)].
- [104] P. Athron, C. Balázs, *et. al.*, *Global fits of axion-like particles to XENON1T and astrophysical data*, *Journal of High Energy Physics* **2021** (2021) 159, [[arXiv:2007.05517](#)].
- [105] P. Carena, O. Straniero, *et. al.*, *Constraints on the coupling with photons of heavy axion-like-particles from Globular Clusters*, *Physics Letters B* **809** (2020) 135709, [[arXiv:2004.08399](#)].
- [106] BOSS: S. Alam *et. al.*, *The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: cosmological analysis of the DR12 galaxy sample*, *MNRAS* **470** (2017) 2617–2652, [[arXiv:1607.03155](#)].
- [107] D. Scolnic *et. al.*, *The Complete Light-curve Sample of Spectroscopically Confirmed SNe Ia from Pan-STARRS1 and Cosmological Constraints from the Combined Pantheon Sample*, *Astrophys. J.* **859** (2018) 101, [[arXiv:1710.00845](#)].
- [108] GAMBIT Scanner Workgroup: G. D. Martinez, J. McKay, *et. al.*, *Comparison of statistical sampling methods with ScannerBit, the GAMBIT scanning module*, *Eur. Phys. J. C* **77** (2017) 761, [[arXiv:1705.07959](#)].
- [109] S. Algeri, J. Aalbers, K. D. Morâ, and J. Conrad, *Searching for new phenomena with profile likelihood ratio tests*, *Nature Reviews Physics* **2** (2020) 245–252, [[arXiv:1911.10237](#)].
- [110] CMB-S4: K. N. Abazajian *et. al.*, *CMB-S4 Science Book, First Edition*, [[arXiv:1610.02743](#)].
- [111] K. Abazajian *et. al.*, *CMB-S4 Science Case, Reference Design, and Project Plan*, [[arXiv:1907.04473](#)].
- [112] W. J. Handley, M. P. Hobson, and A. N. Lasenby, *POLYCHORD: next-generation nested sampling*, *MNRAS* **453** (2015) 4384–4398, [[arXiv:1506.00171](#)].
- [113] R. Trotta, *Bayes in the sky: Bayesian inference and model selection in cosmology*, *Contemp. Phys.* **49** (2008) 71–104, [[arXiv:0803.4089](#)].
- [114] J. Skilling, *Nested sampling for general bayesian computation*, *Bayesian Anal.* **1** (2006) 833–859.
- [115] G. Ashton, N. Bernstein, *et. al.*, *Nested Sampling for physical scientists*, *Nature Reviews Physics (In press)* (2022).
- [116] L. T. Hergt, W. J. Handley, M. P. Hobson, and A. N. Lasenby, *Bayesian evidence for the tensor-to-scalar ratio r and neutrino masses m_ν : Effects of uniform vs logarithmic priors*, *Phys. Rev. D* **103** (2021) 123511, [[arXiv:2102.11511](#)].
- [117] W. Handley, *anesthetic: nested sampling visualisation*, *The Journal of Open Source Software* **4** (2019) 1414, [[arXiv:1905.04768](#)].
- [118] D. Baumann, D. Green, and B. Wallisch, *New Target for Cosmic Axion Searches*, *Phys. Rev. Lett.* **117** (2016) 171301, [[arXiv:1604.08614](#)].
- [119] F. D’Eramo, R. Z. Ferreira, A. Notari, and J. L. Bernal, *Hot axions and the H_0 tension*, *JCAP* **2018** (2018) 014, [[arXiv:1808.07430](#)].
- [120] D. Ghosh and D. Sachdeva, *Constraints on Axion-Lepton coupling from Big Bang Nucleosynthesis*, *JCAP* **2020** (2020) 060, [[arXiv:2007.01873](#)].
- [121] D. Green, Y. Guo, and B. Wallisch, *Cosmological implications of axion-matter couplings*, *JCAP* **2022** (2022) 019, [[arXiv:2109.12088](#)].

- [122] G. Lucente, P. Carenza, T. Fischer, M. Giannotti, and A. Mirizzi, *Heavy axion-like particles and core-collapse supernovae: constraints and impact on the explosion mechanism*, *JCAP* **2020** (2020) 008, [[arXiv:2008.04918](#)].
- [123] F. Calore, P. Carenza, M. Giannotti, J. Jaeckel, and A. Mirizzi, *Bounds on axionlike particles from the diffuse supernova flux*, *Phys. Rev. D* **102** (2020) 123005, [[arXiv:2008.11741](#)].
- [124] A. Caputo, G. Raffelt, and E. Vitagliano, *Muonic boson limits: Supernova redux*, *Phys. Rev. D* **105** (2022) 035022, [[arXiv:2109.03244](#)].
- [125] A. Goudelis, M. Pospelov, and J. Pradler, *Light Particle Solution to the Cosmic Lithium Problem*, *Phys. Rev. Lett.* **116** (2016) 211303, [[arXiv:1510.08858](#)].
- [126] P. Scott, *Pippi – painless parsing, post-processing and plotting of posterior and likelihood samples*, *Eur. Phys. J. Plus* **127** (2012) 138, [[arXiv:1206.2245](#)].
- [127] GAMBIT Cosmology Workgroup: P. Stöcker *et. al.*, *Strengthening the bound on the mass of the lightest neutrino with terrestrial and cosmological experiments*, *Phys. Rev. D* **103** (2021) 123508, [[arXiv:2009.03287](#)].
- [128] BICEP2 Collaboration, Keck Array Collaboration, *et. al.*, *Constraints on Primordial Gravitational Waves Using Planck, WMAP, and New BICEP2/Keck Observations through the 2015 Season*, *Phys. Rev. Lett.* **121** (2018) 221301, [[arXiv:1810.05216](#)].
- [129] GAMBIT Collider Workgroup: C. Balázs, A. Buckley, *et. al.*, *ColliderBit: a GAMBIT module for the calculation of high-energy collider observables and likelihoods*, *Eur. Phys. J. C* **77** (2017) 795, [[arXiv:1705.07919](#)].
- [130] GAMBIT Flavour Workgroup: F. U. Bernlochner, M. Chrzaszcz, *et. al.*, *FlavBit: A GAMBIT module for computing flavour observables and likelihoods*, *Eur. Phys. J. C* **77** (2017) 786, [[arXiv:1705.07933](#)].
- [131] M. Chrzaszcz, M. Drewes, *et. al.*, *A frequentist analysis of three right-handed neutrinos with GAMBIT*, *Eur. Phys. J. C* **80** (2020) 569, [[arXiv:1908.02302](#)].
- [132] GAMBIT Models Workgroup: P. Athron, C. Balázs, *et. al.*, *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](#)].
- [133] S. Bloor, T. E. Gonzalo, *et. al.*, *The GAMBIT Universal Model Machine: from Lagrangians to likelihoods*, *European Physical Journal C* **81** (2021) 1103, [[arXiv:2107.00030](#)].
- [134] M. Giannotti, L. D. Duffy, and R. Nita, *New constraints for heavy axion-like particles from supernovae*, *JCAP* **2011** (2011) 015, [[arXiv:1009.5714](#)].
- [135] A. Caputo, P. Carenza, *et. al.*, *Axionlike Particles from Hypernovae*, *Phys. Rev. Lett.* **127** (2021) 181102, [[arXiv:2104.05727](#)].
- [136] L. Oberauer, C. Hagner, G. Raffelt, and E. Rieger, *Supernova bounds on neutrino radiative decays*, *Astroparticle Physics* **1** (1993) 377–386.
- [137] GAMBIT Collaboration: P. Athron, C. Balázs, *et. al.*, *GAMBIT: The Global and Modular Beyond-the-Standard-Model Inference Tool. Addendum for GAMBIT 1.1: Mathematica backends, SUSYHD interface and updated likelihoods*, *Eur. Phys. J. C* **78** (2018) 98, [[arXiv:1705.07908](#)]. Addendum to [32].