Constraints on axionlike particles from a combined analysis of three flaring *Fermi* flat-spectrum radio quasars

James Davies[1,*] Manuel Meyer,[2,†,§] and Garret Cotter[1,‡]

[1]*University of Oxford, Department of Physics, Oxford OX1 3PJ, United Kingdom*
[2]*Institute for Experimental Physics, University of Hamburg,*
Luruper Chaussee 149, D-22761 Hamburg, Germany

Many theories beyond the Standard Model of particle physics predict the existence of axionlike particles (ALPs) that mix with photons in the presence of a magnetic field. Searching for the effects of ALP-photon mixing in gamma-ray observations of blazars has provided some of the strongest constraints on ALP parameter space so far. Previously, only individual sources have been analyzed. We perform a combined analysis on *Fermi* Large Area Telescope data of three bright flaring flat-spectrum radio quasars, with the blazar jets themselves as the dominant mixing region. For the first time, we include a full treatment of photon-photon dispersion within the jet and account for the uncertainty in our B-field model by leaving the field strength free in the fitting. Overall, we find no evidence for ALPs but are able to exclude the ALP parameters $5 \text{ neV} \lesssim m_a \lesssim 200 \text{ neV}$ and $g_{a\gamma} \gtrsim 5 \times 10^{-12} \text{ GeV}^{-1}$ with 95% confidence.

I. INTRODUCTION

Axions are very light pseudoscalar particles beyond the Standard Model (SM) of particle physics [1–3], which provide a theoretical solution to the strong *CP* problem [4]. Importantly, axions would couple to photons in the presence of an external magnetic field, with a coupling $g_{a\gamma}$, proportional to its mass m_a [5,6]. This coupling would lead to oscillations between photons and axions, comparable to those between neutrino states—an effect that has been the basis for many experimental axion searches (e.g., [7]). So far, none have been found.

Axionlike particles (ALPs) are similar particles in which the $m_a/g_{a\gamma}$ relation is relaxed. Such particles commonly arise in string theories, or as pseudo-Nambu-Goldstone bosons in other SM extensions [8–11]. ALPs would no longer necessarily solve the strong *CP* problem, but they are good candidates to make up all or some of the dark matter content of the Universe [12–15]. This makes them interesting targets for direct and indirect searches too

(e.g., [7,16,17]). In particular, ALP-photon mixing in the various magnetic fields found in space could affect observations of astrophysical sources (e.g., [18,19]). X- and gamma-ray observations of blazars have been used to set some of the strongest constraints on ALP parameter space so far for masses $m_a \lesssim 100 \text{ neV}$ [20–25].

Blazars are active galactic nuclei (AGN) producing jets of relativistic plasma, which are pointed towards us (within a few degrees). This means their emission is strongly enhanced by relativistic effects; blazars make up some of the brightest gamma-ray sources in the sky [26], though they emit across the entire electromagnetic spectrum, from radio to gamma rays. While the detailed emission mechanisms of blazars are still unclear, the low energy emission is usually considered to be synchrotron photons emitted by electrons in the plasma. The high energy emission is then thought to be inverse-Compton (IC) emission from these same electrons up-scattering either their own synchrotron photons (synchrotron self-Compton), or other background photon fields (external Compton) [27]. Hadronic models are also possible for the high energy peak, in particular, (e.g., [28,29]), though these models may require super-Eddington jets [30]. Significantly for ALP searches, a smooth nonthermal distribution of electrons (as produced by, e.g., shock acceleration [31,32]) would produce intrinsically smooth gamma-ray spectra. The presence of ALPs, however, could produce oscillatory spectral features, as the ALP-photon oscillation length could be energy-dependent for some astrophysical environments along the line of sight to the source [33]. Looking for these irregularities in individual blazar spectra

*[*]james.davies2@physics.ox.ac.uk
[†]manuel.meyer@desy.de
[‡]garret.cotter@physics.ox.ac.uk
[§]Also at CP3-Origins, University of Southern Denmark, Campusvej 55, 5230 Odense M, Denmark.

TABLE I.   Time ranges (t_{start} to t_{end}) and best-fit spectral parameters [see Eqs. (1) and (2)] for a combined event-type analysis of the flares. N_0 is the spectral normalization, E_0 is the reference energy, E_c is the cutoff energy, and Γ_1, Γ_2 and κ are indices.

t_{start} MJD	t_{end} MJD	N_0 [10^{-9} MeV^{-1} cm^{-2} s^{-1}]	Γ_1	κ	E_c MeV	Γ_2	E_0 MeV
			3C454.3 (4FGL J2253.9 + 1609)				
55,516.55	55,525.48	525.7 ± 48.46	1.443 ± 0.029		2.614 ± 0.414	0.227 ± 0.0046	410.0
			CTA 102 (4FGL J2232.6 + 1143)				
57,749.10	57,754.09	2.113 ± 0.175	1.813 ± 0.036		9848 ± 2315	0.819 ± 0.134	1000
			3C279 (4FGL J1256.1 − 0547)				
57,188.07	57,189.94	12.47 ± 0.262	2.004 ± 0.018	0.126 ± 0.013			442.1

(NGC 1275 and PKS 2155-304), using their magnetized cluster environments as the mixing region, has been the basis for constraints with *Fermi* Large Area Telescope (LAT), High Energy Stereoscopic System, and Chandra observations [20–22]. These searches require good statistics in the gamma-ray data, which is why bright blazars make good targets—especially when in a flaring state.

Here, we perform a similar search with a combined analysis of *Fermi*-LAT data for three bright flaring flat-spectrum radio quasars (FSRQs; 3C454.3, CTA 102, and 3C279), using the blazar jets themselves as the main mixing regions. It has been suggested that the strong field in the jet could lead to ALP-photon mixing at higher masses than previously probed by gamma-ray searches [34–40], though so far no search has been performed using it as a mixing region. By combining observations from multiple sources, it should also be possible to strengthen the constraints—within the parameter space probed by all the sources—if an ALP signature is seen in one source, then it should be seen in the others too.

In Sec. II we outline the data selection and spectral analysis performed on the three sources. Then, in Sec. III we describe the jet and photon-field modeling required to calculate the ALP-photon oscillations produced within the sources. In Sec. IV, we then discuss the fitting and statistical analysis used to compare the ALP and no-ALP hypotheses and place limits on ALP parameter space, before presenting the results of our analysis in Sec. V. Details of the field structure parameters and the spectral energy distribution (SED) modeling are discussed further in Appendices A and B, respectively, and the effects of systematics are discussed in Appendix C.

II. DATA ANALYSIS

The LAT is a pair-conversion imaging gamma-ray detector on board the *Fermi Gamma-ray Space Telescope* (*Fermi*), which measures gamma rays from 30 MeV up to > 300 GeV energies [41]. Our aim is to look for oscillations in *Fermi*-LAT gamma-ray spectra caused by ALP-photon mixing. We target the three sources

with the brightest flares over the *Fermi* lifetime: 3C454.3, CTA 102, and 3C279 [42]. We use FERMIPY v1.0.1[1] [43] and *Fermi Science Tools* v2.0.8[2] for the analysis.

A. Data selection

Initially, we analyze each source over a significant fraction of the *Fermi*-LAT lifetime (11.7 years between August 4, 2008 and April 1, 2020) to get an average model for each region of interest (ROI). We choose an energy range of 100 MeV to 500 GeV. This long-term ROI model can then be used as an initial condition for fitting the flare observations. Each ROI is centered on the respective source and has a size of $15° \times 15°$. To avoid including gamma-rays produced from the Earth limb, we only use events with a zenith angle $\theta_z \leq 90°$. We choose a spatial binning of 0.1° pixel^{-1}. We use the P8R3_SOURCE_V2 instrument response functions[3] (IRFs) and only use events that pass the P8R3 SOURCE event selection. Because we are looking for spectral oscillations, we make use of the EDISP event classes available with the Pass 8 IRFs [44]. Events are classified into four classes, EDISP0 to EDISP3, depending on the quality of their energy reconstruction (worst to best, respectively). These classes each contain a similar number of events and are analyzed separately with their corresponding IRFs. This allows us to extract the best spectral information from the data possible. For the long-term analysis, we use eight energy bins per decade. Then, for the flare analyses, we choose the binning so as to reach the smallest resolvable energy scale. This is done by extracting the detector response matrices for our observations and choosing the bin width to match the minimum $\Delta E/E$ for the best energy dispersion class (EDISP3).

This gives 65 (3C454.3), 67 (CTA 102), and 61 (3C279) bins per decade for our sources.

B. ROI fitting

First, we optimize the ROI model for each of our sources, for all the event types combined, over the entire 11.7 year time range defined above. The initial model includes every point source in the 4FGL catalogue (Data Release 1) [26] and the standard diffuse isotropic and galactic background templates.[4] We then free the normalization of all sources, including the diffuse backgrounds. Point sources within 5° of the ROI center or with test statistic TS > 10 have the rest of their spectral parameters freed too (TS is the log-likelihood ratio of the likelihoods with and without the source). We free the spectral index of the galactic background as well. These free model parameters are then fitted to the data. Within this fitted ROI, we search for new point sources to add to the model by calculating a TS map. This is done by adding a potential point source (with a power-law index, $\Gamma = 2$) at each pixel of the ROI and calculating its TS. Sources with $\sqrt{TS} \geq 5$ are then added to the overall ROI model at the position which gives the highest TS. We then reoptimize the entire ROI and repeat the process until no more sources are found; overall we find four new sources each for 3C454.3 and 3C279, and three for CTA 102. This gives the final best-fit model for each of our ROIs over the long-term time period. The time ranges used for our flares are taken from the light-curve analysis of [42] and are listed in Table I. For each of these ranges, we redo the above analysis using these final best-fit ROI models, including the new sources, as the initial conditions—this time only freeing the galactic background and sources that still have TS > 10. We fit each event type separately, treating them as separate measurements (as in [21]). Once fitted ROI models have been found for each of our flares, we calculate SEDs for our sources of interest. Following the 4FGL catalogue, the spectra of 3C454.3 and CTA 102 are both best fitted by a power law with a superexponential cutoff:

$$\frac{dN}{dE} = N_0 \left(\frac{E}{E_0}\right)^{-\Gamma_1} \exp\left\{-\left(\frac{E}{E_c}\right)^{\Gamma_2}\right\}, \quad (1)$$

whereas the 3C279 spectrum is best fitted with a log-parabola,

$$\frac{dN}{dE} = N_0 \left(\frac{E}{E_0}\right)^{-(\Gamma_1 + \kappa \ln(E/E_0))}. \quad (2)$$

---

[4]We use `iso_P8R3_SOURCE_V2` templates for the isotropic background for each `EDISP` class and `gll_iem_v07.fits` for the galactic background, which can be found here: https://fermi.gsfc.nasa.gov/ssc/data/access/lat/BackgroundModels.html as accessed on October 5, 2022.
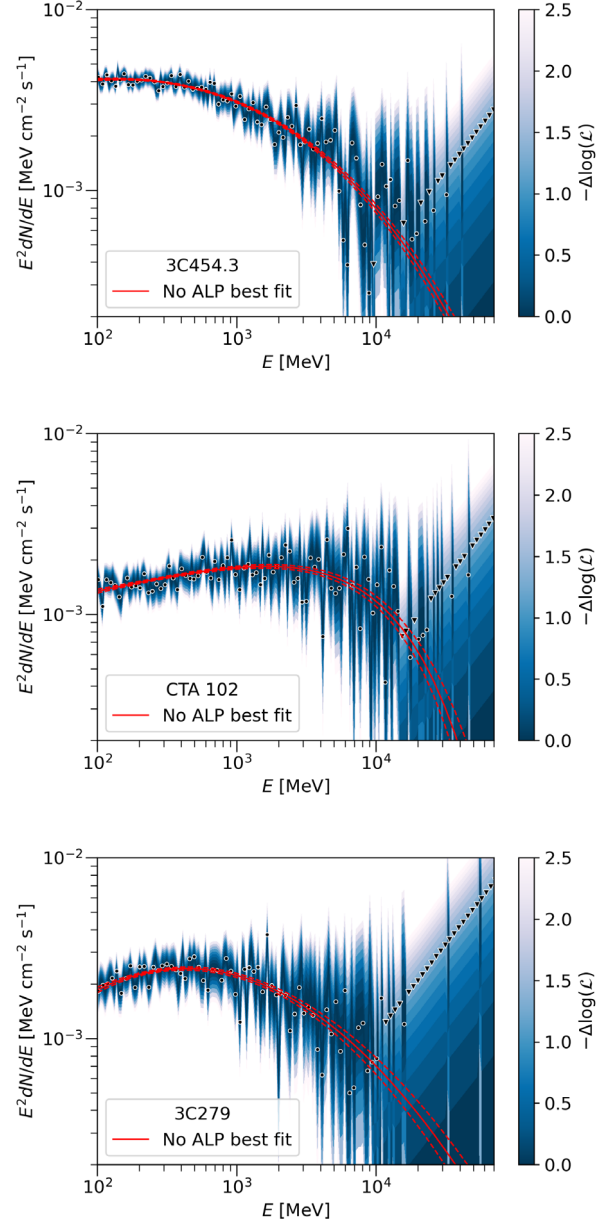

FIG. 1. SEDs for our sources during the flares: 3C454.3 (top), CTA 102 (middle), and 3C279 (bottom). Best-fit spectral parameters (corresponding to the red lines) and time ranges used are listed in Table I. Black points show detections, and triangles show 95% upper limits. Shaded regions show the likelihood curves for each energy bin. For clarity, only points for every other energy bin are plotted.

N is the number of photons received per unit area per unit time at photon energy E, N_0 is the spectral normalization, E_0 is the reference energy, E_c is the cutoff energy, and Γ_1, Γ_2, and κ are indices. Each event type will have different best-fit spectral parameters; those for a combined event-type analysis are shown in Table I, and the corresponding SEDs ($E^2 dN/dE$) are shown in Fig. 1. For clarity, only every other energy bin is plotted, but we utilize the full energy resolution in our analysis steps.

For each event type k, we then extract likelihood curves,[5] $\mathcal{L}^k(\mu_i)$, in each energy bin i, as a function of expected counts μ_i, from these best-fit SEDs[6] (shown as blue bands in Fig. 1). As can be seen, the best statistics are at low energies, and no detected emission is seen at energies above about 80 GeV. These curves can then be used to perform a likelihood ratio test between models with and without ALPs (see Sec. IV). For each event type, the total likelihood for the no-ALPs model is then

$$\mathcal{L}_0^k = \prod_i \mathcal{L}^k(\bar{\mu}_i) \qquad (3)$$

for each source, where $\bar{\mu}_i$ are the expected counts from the best-fit spectral models, including all photon absorption (see Fig. 3 below), but without ALPs (best-fits with ALPs will later be denoted with a hat as opposed to a bar).

III. ALP-PHOTON OSCILLATIONS

To test whether an ALP signature is present in the data, we must model the spectral oscillations caused by ALP-photon mixing.

In general, a photon of energy E propagating in a homogeneous field B (with a component B_T, transverse to the photon direction of motion, and parallel to one of the photon polarization states) will oscillate into an ALP with mass m_a and coupling $g_{a\gamma}$, with a wave number [33,45,46]

$$\Delta_{\mathrm{osc}} = \left[\left\{ \frac{m_a^2 - m_T^2}{2E} + E\left(b + \chi + i\frac{\Gamma_{\gamma\gamma}}{2E}\right) \right\}^2 + (g_{a\gamma}B_T)^2 \right]^{\frac{1}{2}}, \qquad (4)$$

where m_T is the effective mass of the photon (see Ref. [40] for the calculation within the jet). χ and $\Gamma_{\gamma\gamma}$ are the total dispersion and absorption terms for the surrounding photon fields, respectively, and

$$b = \frac{7\alpha}{90\pi}\left(\frac{B_T}{B_{\mathrm{cr}}}\right)^2 \qquad (5)$$

is the vacuum QED term describing dispersion off the magnetic field, with B_{cr} the critical magnetic field $B_{cr} = m_e^2/|e| \sim 4.4 \times 10^{13}$ G, where e is the electric charge. Assuming absorption is small, this means there are two so-called "critical energies," around which the oscillation length depends strongly on energy (and so ALP-photon mixing could lead to oscillations in energy spectra):

$$E_{\mathrm{crit}}^{\mathrm{low}} = \frac{|m_a^2 - m_T^2|}{2g_{a\gamma}B_T}, \qquad (6)$$

which depends on the effective mass difference between the ALP and the photon, and

$$E_{\mathrm{crit}}^{\mathrm{high}} = \frac{g_{a\gamma}B_T}{b + \chi}, \qquad (7)$$

which depends on the dispersion terms. For astrophysical plasma environments, these energies can be in the gamma-ray energy range for interesting ALP parameters. This has been the basis of previous searches and is the basis of ours.

In order to model these spectral oscillations, then, we need a model of the field strength and orientation of the magnetic fields along the sight between us and the gamma-ray sources. This allows us to calculate the photon survival probability $P_{\gamma\gamma}$, i.e., the probability that the emitted photon arrives as a photon at Earth as a function of photon energy (taking into account both photon-ALP conversion and absorption via pair production). The magnetic fields we include along the line of sight are the jet field and the galactic magnetic field (GMF) of the Milky Way, as we choose a mass range where the intergalactic magnetic field does not contribute strongly (see Sec. IV below), and these sources are not thought to be in highly magnetized clusters. For the GMF, we use the model of Ref. [47], as used in, e.g., [21,48]. We also include extragalactic background light (EBL) absorption for propagation through intergalactic space, using the model of Ref. [49]. The dominant mixing region we are using, however, is the jet field. We use the gammaALPs Python package[7] to solve the ALP-photon mixing equations throughout—see [52] for an overview.

A. Jet modeling

For mixing within a jet, the detailed structure of the jet field needs to be taken into account, as well as dispersion and absorption from the various photon fields within the jet [39,40,53].

We use the Potter and Cotter jet framework (PC, see [54]) for the overall jet properties (shape of the field strength, bulk Lorentz factors, and electron density)[8] as discussed in the context of ALP-photon mixing in [40]. The structure of the PC jet model is a parabolic magnetically dominated accelerating jet base, which transitions to a decelerating ballistic conical jet in rough energy equipartition at $r_{\mathrm{tr}} \sim 10^5 r_g$ from the black hole, where r_g is the gravitational radius, which depends on the black hole mass as $r_g = 2GM/c^2$. In the accelerating region ($r \leq r_{\mathrm{tr}}$) the bulk Lorentz factor is $\Gamma \propto r^{1/2}$, and it is $\Gamma \propto \log(r)$ in the

---

[5]Throughout, we use the shorthand $\mathcal{L}(\mu) \equiv \mathcal{L}(\mu|x)$, where x is the observed data.

[6]We extract bin-by-bin likelihood curves using the SED function within FERMIPY. This function changes the normalization in each energy bin and recomputes the likelihood at each point, taking energy dispersion into account.

[7]Hosted on GitHub (https://github.com/me-manu/gammaALPs) and archived on Zenodo [50]. Data files and an example notebook connected to this publication are also available at [51].

[8]Our jet model is available within the gammaALPs package.

TABLE II. Jet properties for our sources: r_{em} and r_{tr} are the locations of the emission region and jet-base transition region, respectively; r_{jet} is the jet length; B_0 is the field strength, n_e is the electron density; Γ is the bulk Lorentz factor; and r_T, α, and f are the field structure parameters.

| Parameters	Unit	3C454.3	CTA 102	3C279
r_{em}	pc	0.103	0.104	0.016
r_{tr}	pc	59.8	56.6	47.9
r_{jet}	kpc	100	75.3	32.4
$B_0(r_{tr})$	mG	16.0	26.2	6.28
$n_e(r_{tr})$	cm^{-3}	4.7	2.5	5.0
$\Gamma(r_{tr})$		60	52	37
$\Gamma(r_{jet})$		35	29	18
r_T	pc	59.8	56.6	47.9
α		1	1	1
f		0.3	0.3	0.3

FIG. 2. One realization of the transverse component of the magnetic field B_T for 3C454.3, using the parameters $f = 0.3$, $\alpha = 1$, and $r_T = r_{tr}$. Vertical dashed line shows r_{tr}.

decelerating region ($r_{tr} < r \leq r_{jet}$), where r_{jet} is the jet length. This leptonic jet framework is consistent with theory, observation, and simulations and is capable of reproducing broadband steady-state SEDs for many blazars [54].

For the location of the gamma-ray emitting regions during the flares r_{em}, we use the lower limits found in Ref. [42], derived from the absence of attenuation due to pair production with broad line region (BLR) photons in the gamma-ray spectra. We use the $B(1 \text{ pc})$ values found in [54] from fits with the PC model to set the initial value of the magnetic field strength B_0, which is then left free in the fitting (see Sec. IV below). These initial values are slightly lower than those derived from very-long-baseline interferometry core-shift measurements for each of our sources in Ref. [55], where they assume a conical jet throughout.[9] The electron density varies as $n_e \propto R^{-2}$, where R is the jet width, with the value at r_{tr} derived from energetic equipartition. Values for the jet parameters used are listed in Table II.

We then model the detailed field structure as in Ref. [40], with a tangled component (B_t) and a helical component (B_h) that transitions from poloidal to toroidal as r increases down the jet. A constant fraction f of the total field energy density is in the tangled component:

$$\frac{B_t^2}{B_h^2} = \frac{f}{1-f}. \tag{8}$$

The radius at which the helical field component becomes toroidal is r_T; the transverse component of the helical field varies as $B_T \propto r^{-\alpha}$ for $r < r_T$. The three parameters f, r_T, and α therefore govern the detailed field structure (along with a treatment of the coherence length of the tangled field). Ideally, these parameters would be allowed to vary in

---

[9]We also found that these larger values were incompatible with the SED modeling performed in Appendix B.
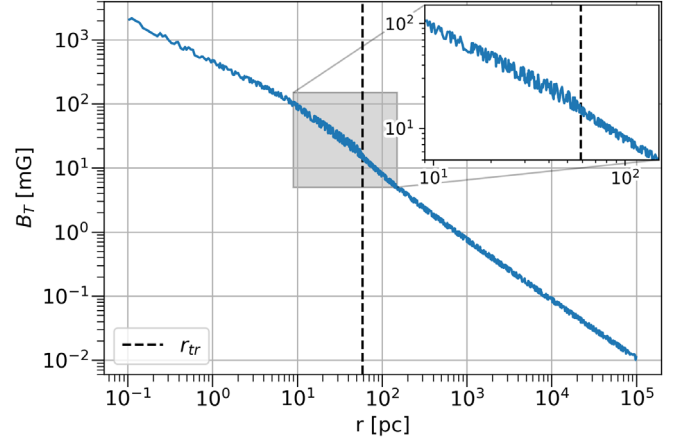
the fit in the same way B_0 is. However, because of computational constraints ($P_{\gamma\gamma}$ has to be recalculated every time one of them changes), it is necessary to fix them. In Appendix A, we motivate our choices for these parameters from observation and simulations and show that varying them would be unlikely to strongly affect our final results. For all our sources, we use $\alpha = 1$ and $r_T = r_{tr}$. Figure 2 shows one example field realization for 3C454.3 with this set of parameters.

B. Photon fields

As well as the magnetic field structure, background photon fields can also affect ALP-photon mixing [46,53] [see χ and $\Gamma_{\gamma\gamma}$ in Eqs. (4) and (7)]. This is because the oscillations are sensitive to slight differences in propagation between the ALP and photon states. Specifically, gamma rays will be affected by photon-photon dispersion and absorption via pair production from background photon fields, whereas ALPs will not. The fields we would expect within FSRQs are those from the central AGN (accretion disk, BLR, dust torus), starlight (extragalactic and from the host galaxy), the cosmic microwave background (CMB), and synchrotron photons from the jet plasma itself. Reference [53] investigated the effects of all these fields on mixing within 3C454.3. They found that for emission regions on the scale of the AGN fields, dispersion off of them will dominate and should be included in the calculations. In particular, for our $r_{em} \sim 0.1$ pc, we expect the BLR and torus fields to be the most important, as the disk is only relevant at much smaller scales. Dispersion from the CMB can play a large role within the jet at energies above 100 GeV, but, as can be seen from Fig. 1, we are only interested in lower energies. (In fact, gamma-rays at these energies would likely be absorbed by BLR photons in our sources anyway, see Fig. 3 below.) Modeling of the starlight and synchrotron fields therefore does not have
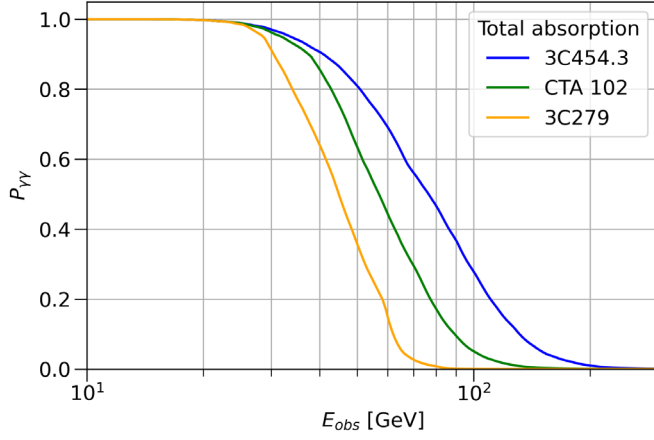
FIG. 3. Photon survival probability $P_{\gamma\gamma}$ as a function of observed energy for each of our sources, displaying the total absorption for each of our sources, including all photon fields.

to be extremely precise. Nevertheless, we include all the photon fields for each of our sources, using the same method and models as Ref. [53]. The various parameters we use for the AGN fields, along with their sources, are given in Table III.

The χ and Γ calculations depend on the geometry as well as the photon energies and energy densities of the background fields. The disk is modeled as flat, extending radially in the plane perpendicular to the jet between R_{in} and R_{out}, with each radius between the two emitting at only

one energy (as in Ref. [56]). We use $R_{\text{in}} = 6r_g$ for all sources, the expected inner disk radius for a Schwarzschild black hole.

The BLR is modeled as a series of concentric rings, each corresponding to an emission line, and also emitting at only one energy. The radii and luminosities of the lines can be derived from those of the Hβ line for each source (we use all the lines in the Appendix of [56]).

We use the torus model described in Ref. [53]—an extension of the flat model of Ref. [56] to include an elliptical torus cross section. Each torus emits at a single energy, depending on its temperature Θ, and the fraction of disk radiation reemitted in each case is assumed to be $\xi_{dt} = 0.1$ (as in [56]). All tori for our sources are given the same size and shape. They extend radially between R_1 and R_2 and have a height so as to give a covering fraction—the fraction of the sky obscured by the torus from the point of view of the black hole—of $f_c = 0.6$, which is considered typical (e.g., [58]). The cloud number density within the torus decreases with R from the black hole $\propto R^{-1}$ for all our sources (see [53,56] for details).

We also include the EBL, starlight, and CMB fields exactly as described in Ref. [53], though, as mentioned above, they are subdominant. In order to model the (also subdominant) synchrotron photon field within the jet, we then follow Ref. [53] in modeling broadband SEDs for each of our sources (in both flaring and steady states) and compare them to observations (see Appendix B). This also enables us to confirm the overall self-consistency of our jet and photon-field models.

TABLE III. Parameters used for the AGN fields. From the top: disk luminosity, black hole mass (for setting the gravitational radius r_g), inner disk radius, outer disk radius, Hβ line luminosity, Hβ line radius, torus temperature, inner torus radius, outer torus radius, semiminor to semimajor axis ratio of torus cross section.

Parameter	Unit	3C454.3[a]	CTA 102[b]	3C279[c]
		Disk		
L_{disk}	erg s^{-1}	2×10^{46}	4×10^{46}	3×10^{45}
M_{BH}	$M_\odot$	1.2×10^9	8.51×10^8	3×10^8
R_{in}	r_g	6	6	6
R_{out}	r_g	200	200[a]	430
		Broad Line Region		
$L_{H\beta}$	erg s^{-1}	4.18×10^{43}	4.93×10^{43}	$1.73 \times 10^{43\,\text{b}}$
$R_{H\beta}$	cm	4.3×10^{17}	6.1×10^{17}	2.8×10^{17}
		Torus		
Θ	K	1000	1000	500
R_1	cm	1.6×10^{19}	$1.6 \times 10^{19\,\text{a}}$	$1.6 \times 10^{19\,\text{a}}$
R_2	cm	1.6×10^{20}	$1.6 \times 10^{20\,\text{a}}$	$1.6 \times 10^{20\,\text{a}}$
b/a[d]		0.527 ($f_c = 0.6$)	0.527 ($f_c = 0.6$)	0.527 ($f_c = 0.6$)

[a]Values taken from Ref. [56], unless marked otherwise.
[b]From Ref. [42] unless marked otherwise.
[c]From Ref. [57] unless marked otherwise. For $R_{H\beta}$ for 3C279 we use the relation of $R_{H\beta}$ and $R_{Ly\alpha}$ in Ref. [56] to convert from $R_{Ly\alpha}$.
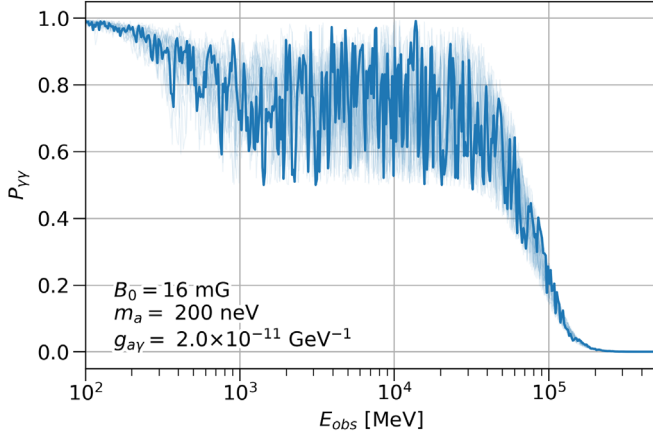[d]From Ref. [53].

FIG. 4. Example photon survival probability $P_{\gamma\gamma}$ for one realization of the 3C454.3 field, using $f = 0.3$, $\alpha = 1$, and $r_T = r_{\mathrm{tr}}$. Twenty more realizations are shown in the background. Dispersion and absorption off of the background photon fields are included.

Figure 3 shows photon survival probability $P_{\gamma\gamma}$ as a function of observed energy for each of our sources, displaying the total absorption from all the photon fields (including intergalactic EBL absorption). The absorption rates are calculated as in Ref. [53] and are included in every calculation, both with and without ALPs. We note that, even though B_0 changing in the fit would, in principle, change the synchrotron field, we keep all the fields fixed. This is a good approximation because the synchrotron field hardly affects the dispersion and never affects the absorption (see [53] and Appendix B).

IV. STATISTICAL METHODS

We are now in a position to compute photon survival probabilities, $P_{\gamma\gamma}(m_a, g_{a\gamma}, \mathbf{B}_j)$, for ALP-photon beams propagated through our sources, where $\mathbf{B} = (B_0, f, \alpha, r_T)$ and j denotes a realization of the random magnetic field. Figure 4 shows an example $P_{\gamma\gamma}(E_{\mathrm{obs}})$ for one pair of ALP parameters for 3C454.3, including both dispersion and absorption from the background fields. Our aim is to compare models with ALPs to the observed *Fermi* data. We follow the statistical methods of, e.g., Refs. [20,21,48] closely. For a random B-field realization j and spectral parameters $\boldsymbol{\theta}$, the expected counts including ALPs are then

$$\mu_i(m_a, g_{a\gamma}, j) = \langle P_{\gamma\gamma}(m_a, g_{a\gamma}, \mathbf{B}_j) \rangle_i \cdot \mu_i(\boldsymbol{\theta}), \quad (9)$$

where $\langle P_{\gamma\gamma} \rangle_i$ denotes the average over energy bin i, which is necessary because $P_{\gamma\gamma}$ can vary on energy scales much smaller than the bin width. It is worth noting that possible uncertainties in the instrument response functions used within FERMIPY could possibly slightly affect this expression. In Appendix C, we show that the inclusion of an extra shift or smear in the energy reconstruction or dispersion would not

greatly affect our overall results. We calculate $P_{\gamma\gamma}$ at 500 fine energy bins, logarithmically spaced across our energy range, before averaging. For each source, one set $(m_a, g_{a\gamma}, \mathbf{B}_j, \boldsymbol{\theta})$ corresponds to a likelihood,

$$\mathcal{L}_{\mathrm{ALP}}(m_a, g_{a\gamma}, \mathbf{B}_j, \boldsymbol{\theta}) = p(B_0)\prod_i \mathcal{L}(\mu_i(m_a, g_{a\gamma}, j)), \quad (10)$$

where $\mathcal{L}(\mu_i)$ are the likelihood curves extracted from the *Fermi* SEDs, evaluated at the expected ALP counts, and

$$p(B_0) = \exp\left\{-\frac{1}{2}\left(\frac{B_0 - \bar{B}_0}{\sigma_B}\right)^2\right\} \quad (11)$$

is the prior on B_0, which takes the form of a Gaussian. $\bar{B}_0$ is the initial value used (see Table III), and σ_B is the error derived for the magnetic field strength in Ref. [55].[10] In each case, σ_B is around 20% of $\bar{B}_0$. For each field realization, we fit the ALP spectrum to the data by varying B_0 and $\boldsymbol{\theta}$ in such a way as to maximize $\mathcal{L}_{\mathrm{ALP}}(m_a, g_{a\gamma}, \mathbf{B}_j, \boldsymbol{\theta})$. We use the IMINUIT Python package for the fitting. Note that every time B_0 is changed in the fit, $P_{\gamma\gamma}$ has to be recalculated completely (as is the case when changing m_a or $g_{a\gamma}$). When B_0 changes, only the overall field strength is affected; the random field orientations and domain lengths remain the same for a given j. As shown in Appendix B, changing the synchrotron photon field has a negligible effect on dispersion, so we can keep it constant as well. Large variations of B_0, such as removing the field completely, are discouraged by the prior term in the likelihood. Best-fit values of the field strength and spectral parameters are denoted by $\hat{\mathbf{B}}_j$ and $\hat{\boldsymbol{\theta}}$.

We scan an 8×7 logarithmically spaced grid in $(m_a, g_{a\gamma})$-space, with $m_a \in [5, 5000]$ neV and $g_{a\gamma} \in [0.1, 10] \times 10^{-11}$ GeV^{-1}. This region is where we might expect mixing in the jets (see Ref. [40]), with the critical energies [Eqs. (6) and (7)] lying around the *Fermi* energy range. For masses $m_a < 5$ neV, the precise jet length becomes important, as does conversion in the intergalactic magnetic field. Higher couplings are ruled out by experiment [59], and lower couplings would lead to oscillations too small to be detectable.

In order to treat the random field statistically (following, e.g., Refs. [21,48]), for each $(m_a, g_{a\gamma})$ we perform the fits for 100 magnetic field realizations, then sort them by $\mathcal{L}_{\mathrm{ALP}}$ and choose $j = 95$ corresponding to the 0.95 magnetic field realization quantile. Each point on the ALP grid then corresponds to a likelihood value, $\mathcal{L}_{\mathrm{ALP}}^k(m_a, g_{a\gamma}, \hat{\mathbf{B}}_{95}, \hat{\boldsymbol{\theta}})$, for each event type.

For each source, the overall ALP and no-ALP hypotheses can be compared with the test statistic,

---

[10]Errors are derived from their Eq. (4) and the errors on the values quoted in their Table 1 and references, therein.
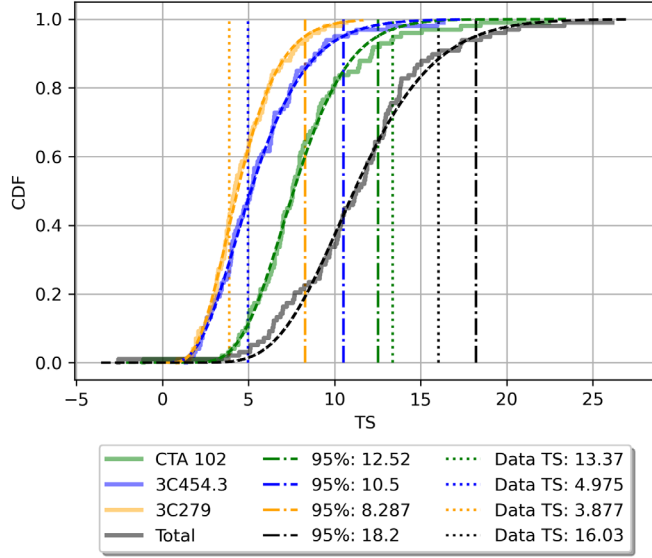
FIG. 5. Cumulative distribution functions (CDF) for the TS values for the individual sources (solid) and the total TS_{tot} distribution (black). Dashed lines show the best-fit gamma distributions. Dot-dashed vertical lines show the 95% thresholds, and dotted vertical lines show the TS values of the data.

$$TS = -2 \sum_{k} \ln \left(\frac{\mathcal{L}_0^k(\bar{\boldsymbol{\theta}})}{\mathcal{L}_{ALP}^k(\hat{m}_a, \hat{g}_{a\gamma}, \hat{\mathbf{B}}_{95}, \hat{\boldsymbol{\theta}})} \right), \quad (12)$$

defined in the standard way for a likelihood ratio test, where $\mathcal{L}_0^k$ are the maximum likelihoods for the no-ALP model, found in Eq. (3), with best-fit spectral parameters $\bar{\theta}$, and the additional hats in the denominator denote the maximum likelihood over the whole ALP grid [60]. A high value of TS would mean that the ALP hypothesis is more likely than the no-ALP hypothesis. To quantify how confidently we could reject the no-ALP hypothesis for a given TS, we use Monte Carlo simulations to find the null TS distribution. For this, we use the `simulate_roi` function in FERMIPY to generate 100 simulated ROIs for each event type (for each of our sources), every time removing the source and injecting a new one from the best-fit spectral model, including photon absorption. The whole analysis is then repeated on each simulated ROI to get a distribution in TS (shown as the solid lines in Fig. 5). The TS thresholds can then be read from these distributions. Specifically, because we only have 100 simulations, we fit gamma distributions to the TS distributions (dashed lines), from which we can take the 0.95 threshold values (dot-dashed lines in Fig. 5). As can be seen, TS > 10.5 is required to reject the no-ALP hypothesis with 95% confidence for 3C454.3; TS > 12.52 for CTA 102; and TS > 8.29 for 3C279.

Regardless of whether we can claim an ALP detection, we would also like to find which ALP parameters a given observation is inconsistent with—i.e., which values of m_a
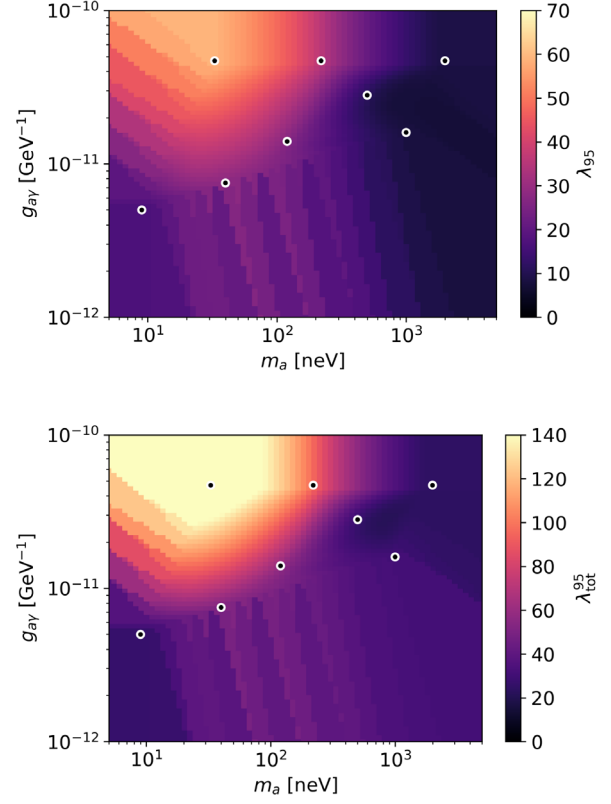

FIG. 6. λ_{95} thresholds for CTA 102 (top) and all sources (λ_{tot}) (bottom). Points show injected $(m_a, g_{a\gamma})$ points, which are interpolated between. Outside the bounds of the injected points, the nearest neighbor is used for the threshold.

and $g_{a\gamma}$ can be excluded. This can be done with a different test statistic:

$$\lambda(m_a, g_{a\gamma}) = -2 \sum_{k} \ln \left(\frac{\mathcal{L}_{ALP}^k(m_a, g_{a\gamma}, \hat{\mathbf{B}}_{95}, \hat{\boldsymbol{\theta}})}{\mathcal{L}_{ALP}^k(\hat{m}_a, \hat{g}_{a\gamma}, \hat{\hat{\mathbf{B}}}_{95}, \hat{\hat{\boldsymbol{\theta}}})} \right), \quad (13)$$

which compares the best fit at each point $(m_a, g_{a\gamma})$ with the overall best fit of the whole grid. A large λ means that the best fit at that point is significantly worse than the best fit overall and so can be rejected by the data. The underlying distribution of $\lambda(m_a, g_{a\gamma})$ can be found in the same way as the TS distribution, except this time an ALP spectrum is injected into the simulated ROIs. This distribution could, in principle, be different for each point in ALP parameter space, as the oscillations do not depend trivially on m_a and $g_{a\gamma}$. Doing the simulations for every point $(m_a, g_{a\gamma})$ is not computationally feasible, however. Therefore, we calculate the λ distribution at various points and linearly interpolate between them to get the 95% λ thresholds across the grid, $\lambda_{95}(m_a, g_{a\gamma})$. For CTA 102, we use eight points, spread over the region of parameter space where we might expect exclusions; the top panel of Fig. 6 shows λ_{95} for CTA 102. The black points show the injected $(m_a, g_{a\gamma})$ pairs. Within

the region bounded by these points, λ_{95} is interpolated, and it can be seen that λ_{95} is not constant but varies smoothly across the grid—generally lowering for decreasing coupling and increasing mass, i.e., for weaker oscillations. Outside the region bounded by the injected points, nearest-neighbor λ_{95} values are used. This is generally a conservative estimate, as λ_{95} would continue to decrease into the unprobed parameter space beyond the lower right edge of the interpolation region, where it would approach the TS threshold. Because of the smoothness of the CTA 102 λ_{95} distribution, we use fewer injected points for 3C454.3 and 3C279—seven and three, respectively—to save on computing time.

So far, we have only discussed individual source analyses. It is possible to combine the likelihoods from the different sources in the same way as those from the different event types within one source (or even different energy bins within one event type). The total ALP and no-ALP likelihoods are just the product of the individual source likelihoods, $\mathfrak{L} = \Pi_s \mathcal{L}_s$, where s indexes the different sources. This means the final $\mathrm{TS}_{\mathrm{tot}}$ and $\lambda_{\mathrm{tot}}(m_a, g_{a\gamma})$ formulae are[11]

$$\mathrm{TS}_{\mathrm{tot}} = -2\sum_k \ln\left(\frac{\mathfrak{L}_0^k(\bar{\boldsymbol{\theta}})}{\mathfrak{L}_{\mathrm{ALP}}^k(\hat{m}_a, \hat{g}_{a\gamma}, \hat{\hat{\mathbf{B}}}_{95}, \hat{\hat{\boldsymbol{\theta}}})}\right), \quad (14)$$

and

$$\lambda_{\mathrm{tot}}(m_a, g_{a\gamma}) = -2\sum_k \ln\left(\frac{\mathfrak{L}_{\mathrm{ALP}}^k(m_a, g_{a\gamma}, \hat{\mathbf{B}}_{95}, \hat{\boldsymbol{\theta}})}{\mathfrak{L}_{\mathrm{ALP}}^k(\hat{m}_a, \hat{g}_{a\gamma}, \hat{\hat{\mathbf{B}}}_{95}, \hat{\hat{\boldsymbol{\theta}}})}\right), \quad (15)$$

where the minima are found after the product over the sources is taken. Of course, because the intrinsic parameters of each source are different, the different sources will be capable of probing slightly different regions of parameter space to greater or lesser degrees. It is important that the likelihoods are combined in this way so that each source contributes proportionally to the overall likelihood. The distributions for these two test statistics can be found in the same way as those for the individual sources. Figure 5 also shows the $\mathrm{TS}_{\mathrm{tot}}$ distribution; a value $\mathrm{TS}_{\mathrm{tot}} > 18.2$ would be required to reject the no-ALP hypothesis with 95% confidence for all the sources combined. The lower panel of Fig. 6 shows the $\lambda_{\mathrm{tot}}^{95}$ thresholds across the ALP parameter space. The same interpolation method is used as before, and again, the overall distribution varies smoothly. For those points where either one or both of 3C454.3 and 3C279 is missing injected simulations, we use the sum of the individual λ_{95} thresholds. This again is a conservative (i.e., over-) estimate of $\lambda_{\mathrm{tot}}^{95}$, as by definition $\lambda_{\mathrm{tot}}^{95} \leq \sum_s \lambda_{95}^s$ everywhere (the two are only equal if the minima

of the likelihood profiles for all the sources lie at the same point). Only the point at 1000 neV does not have injected simulations for both CTA 102 and 3C454.3, which together should dominate the overall thresholds, so in the relevant region of parameter space this approximation is small.

V. RESULTS

Figure 5 also shows (dotted vertical lines) the data TS values for all the sources both individually and in
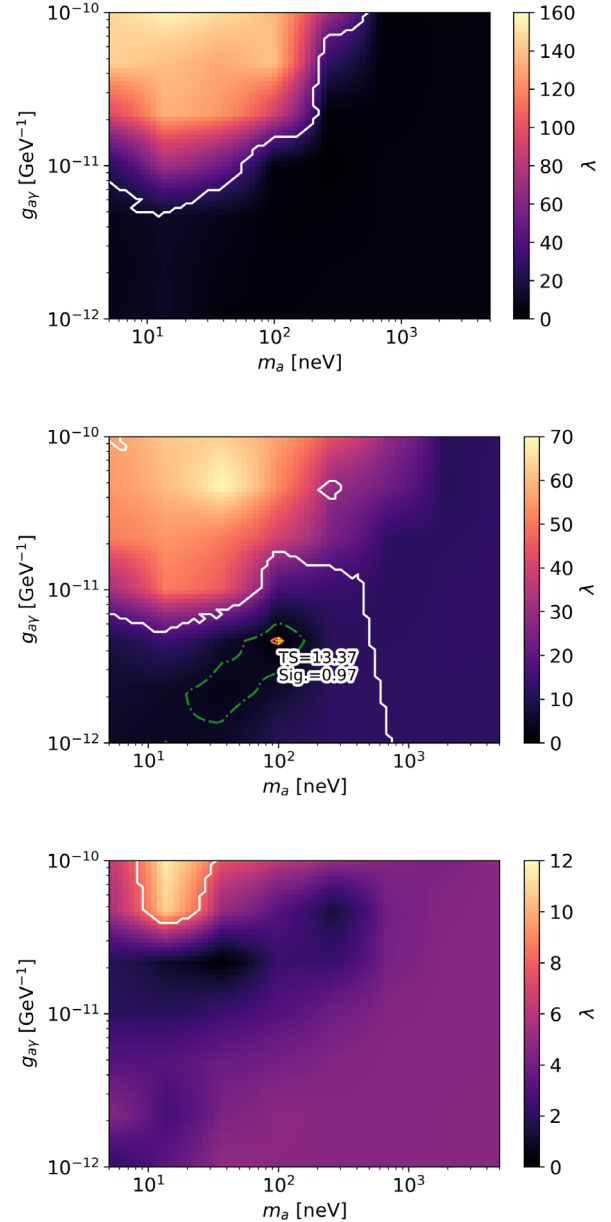

FIG. 7. $\lambda(m_a, g_{a\gamma})$ for 3C454.3 (top), CTA 102 (middle), and 3C279 (bottom). White contours show the 95% exclusions, i.e., they enclose regions where $\lambda \geq \lambda_{95}$. The green dot-dashed and the red solid contours show the 1σ and 2σ preference regions for CTA 102, respectively, and the gold cross shows the location of the best-fit point.

---

[11]Comparable to those used in, e.g., *Fermi* searches for dark matter annihilation lines in dwarf spheroidal galaxies [61] but with the *B*-field taking the role of the *J*-factor in our case.

combination. The TS values for 3C454.3 and 3C279 are below their respective TS thresholds: $TS_{3C454.3} = 4.97$, and $TS_{3C279} = 3.88$. For CTA 102, however, their is a slight preference (2σ) for the ALP case[12]: $TS_{CTA\,102} = 13.37$, which is over the threshold of 12.52. $TS = 13.37$ is in the 97% quantile of the gamma-function fitted to the CTA 102 TS distribution, but falls to the 91% quantile if the distribution is simply read from the simulations. Also, this local significance of $\sim 2\sigma$ for an ALP signal in the CTA 102 data would be further reduced by a trial factor of 3 when considering the fact we looked at three sources. Therefore, this is not a very significant preference for the ALP case, and indeed, it disappears in the combined analysis: $TS_{tot} = 16.03$. Overall then, we cannot rule out the no-ALP hypothesis, or in other words, we have not found an ALP signal in the data.

Nonetheless, we are able to place limits on the parameters m_a and $g_{a\gamma}$. Figure 7 shows λ for each of the sources. The white contours enclose regions, where $\lambda \geq \lambda_{95}$, and so show the 95% exclusion contours for each individual source. For CTA 102, the regions of 1σ (68%) and 2σ (95%) preference over the null hypothesis are also shown as green dot-dashed and red solid contours, respectively. The best-fit point ($m_a = 100.8$ neV and $g_{a\gamma} = 4.64 \times 10^{-12}$ GeV^{-1}) is also plotted as a gold cross, along with its significance (0.97; 2.17σ). As can be seen, because of this slight preference for the ALP case, the 95% exclusions contour from the CTA 102 data extend to high masses and low couplings (which approximates the no-ALP case).

Aside from CTA102, the constraints from 3C454.3 are the strongest, as would be expected from the comparatively good statistics of the 3C454.3 observations (see Fig. 1). Constraints from 3C279 data are much weaker than the other sources not only because its statistics are not quite as good, but also because the configuration of the field parameters means that, with B_0 free, good fits are generally able to be found to the data; for 3C279, λ is smaller for much of the region that is excluded by the other sources than it is in the high-mass–low-coupling region. This highlights the importance of leaving the magnetic field strength free in the fitting.

As was shown in Fig. 5, the preference for the ALP case shown in the CTA 102 data disappears in the combined analysis; we would therefore expect 3C454.3 to contribute most strongly to the combined exclusions. Indeed, Fig. 8 shows λ_{tot} for the whole scanned parameter space, and the 95% exclusions (again shown by the white contour) are only marginally better than the 3C454.3 exclusions. In particular, the high-mass–low-coupling region is not excluded. This highlights the benefits of using a combined analysis to derive robust exclusions. Figure 9 shows how our 95% exclusion contours compare with current

---

[12]For comparison, a 5σ significance is generally required for a new particle detection within the particle physics community.
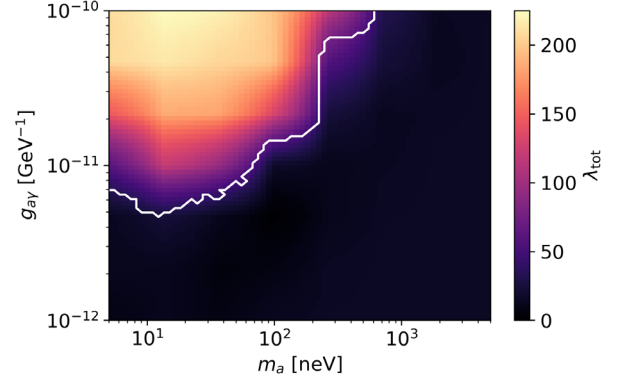

FIG. 8. $\lambda_{tot}(m_a, g_{a\gamma})$ for all the sources combined. White contour shows the 95% exclusions, i.e., enclosing the region where $\lambda_{tot} \geq \lambda_{tot}^{95}$.
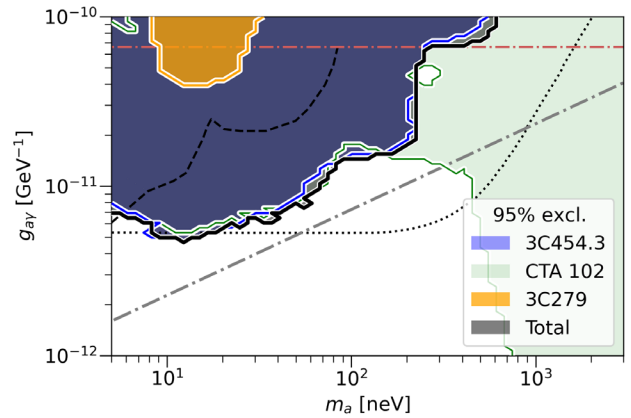

FIG. 9. Overall 95% exclusion contours for each source and for the combined analysis. The black dotted contour shows constraints from magnetic white dwarf radio polarization [62]. Black dashed contours show previous gamma-ray constraints [20,21]. The red dot-dash contour shows the CERN Axion Solar Telescope (CAST) experimental constraints [59]. The dark matter line is shown in grey dot-dash.

constraints, shown by the black and red dashed contours. The dark matter line is shown as a grey dot-dashed line, below which ALPs could make up all of dark matter [15]. As can be seen from the figure, the combined analysis performed here allows the previous gamma-ray constraints to be extended. Overall, we can exclude the parameter space 5 neV $\lesssim m_a \lesssim$ 200 neV and $g_{a\gamma} \gtrsim 5 \times 10^{-12}$ GeV^{-1} with 95% confidence.

VI. CONCLUSIONS

Searches for ALP signals in the high energy spectra of AGN have provided some of the strongest constraints on ALPs m_a and $g_{a\gamma}$ so far [20–22]. These searches have all been analyses of individual sources and have generally used the turbulent magnetic field of the host cluster as their

main mixing region; similar searches are also planned in the future (e.g., [48]).

Here, we have performed, for the first time, a combined analysis on *Fermi*-LAT data of three bright flaring FSRQs (3C454.3, CTA 102, and 3C279), with the blazar jets themselves as the dominant mixing region. These sources were chosen because they displayed the brightest flaring periods over the *Fermi* lifetime.

We analyze each of the sources using the FERMIPY Python package, first over a significant fraction of the *Fermi* lifetime to get average ROI models, which are then used as initial conditions for detailed SED analysis of the flaring time periods. This enables us to extract likelihood curves from the resulting flare SEDs, which can be used to compare ALP spectral models to the data with a log-likelihood ratio test.

To find the ALP spectra, we model the jets within the PC framework, with a helical and a tangled field component as outlined in Ref. [40]. In particular, based on observed polarization fractions of the sources, we use jets with 30% of the magnetic energy density in the tangled component. Also, for the first time, we include a full treatment of photon-photon dispersion within the jet, following Ref. [53]. This requires the modeling of the disk, BLR, torus, synchrotron, starlight, CMB and EBL photon fields within the jets. We have performed SED modeling with our combined jet and photon-field models to ensure they are consistent with both each other and observations.

These jet models then allow us to compute ALP spectra for each of the sources and fit them to the *Fermi* observations. We treat both the tangled field component and the errors in the data statistically, by running the analysis for 100 field simulations on 100 simulated *Fermi* data sets. Also, unlike previous work, we account for the uncertainty in our B-field model by leaving the field strength free in the fits, including a prior term in the likelihood function based on core-shift estimates of the field strength.

To find the underlying distributions of the test statistics used to place limits, we performed the analysis on simulated data with various ALP-spectra injected into it. This was done across the ALP parameter space, enabling a 2D test-statistic threshold to be constructed, as opposed to using a single value everywhere.

In the CTA 102 data, we find a marginal (2σ) preference for ALPs, with the best fit occurring at $m_a = 100.8$ neV and $g_{a\gamma} = 4.64 \times 10^{-12}$ GeV^{-1} (below the dark matter line). This slight preference disappears in the combined analysis, however, highlighting the benefits of using multiple sources. Overall then, we find no evidence for ALPs but are able to exclude the parameter space 5 neV $\lesssim m_a \lesssim 200$ neV and $g_{a\gamma} \gtrsim 5 \times 10^{-12}$ GeV^{-1} with 95% confidence. This is an improvement on previous gamma-ray searches in this mass range, though it is almost completely contained within the magnetic white dwarf polarization

constraints of [62]. Our constraints do not quite reach the dark matter line but are limited in coupling to similar $g_{a\gamma}$ values as previous *Fermi* limits (see [21]), which is to be expected. Nonetheless, we reach lower couplings than those projected to be reached by the future ALPS II experiment in the same mass range [63] and comparable couplings to the projected limits of the future international axion observatory experiment [64].

Future searches with CTA, for instance, (like those outlined in [48]) could likely take advantage of greater instrumental sensitivity to probe lower couplings using this same method, with the blazar jets as the dominant mixing region. It would also be interesting to see how these limits could be extended in the event of another flare as bright as the one from 3C454.3 used here, to fully take advantage of the combined analysis method.

APPENDIX A: FIELD STRUCTURE PARAMETERS

Ideally, all the field parameters (f, α, r_T) would be left free in the fit, in the same way as B_0. Unfortunately, this is not computationally feasible at the moment ($P_{\gamma\gamma}$ has to be recalculated every time a field parameter changes in the fit).

We therefore would like to constrain f, α, and r_T to specific values.

For our sources, at *Fermi* energies, the fraction of magnetic energy density in the tangled component f has the strongest effect on the oscillations. In particular, a very low value of $f < 0.05$ can greatly reduce the magnitude and severity of the oscillations produced by mixing in the jet. Fortunately, f can be somewhat constrained on a source-by-source basis from radio polarization observations. In general, the fractional polarization of radio emission from a source is higher for a uniform field and lower for a disordered field. Reference [65] compare very long baseline array fractional polarization maps of 3C454.3 to simulations, using a uniform helical field. They find that the asymmetry of the maps matches a helical field well, but an additional disordered field component is required to reduce the overall fractional polarization to observed levels. In particular, they need $B_t^2 \sim 0.45 B_h^2$, which equates to $f \sim 0.3$. Fractional polarization maps of CTA 102 (see [66]) look similarly asymmetric and are at similar values. Reference [67] has done similar modeling but with circular polarization as well for 3C279. All of their best-fit models have $0.19 \leq f \leq 0.51$, again ruling out very low tangled field fractions, and their best-fit value is $f = 0.36$. Motivated by these results, we choose a fixed value, $f = 0.3$, for all our sources.

Another factor concerning the tangled field component is which coherence length, l_c, to use. We take the jet width R as an upper limit for the tangled coherence length at a given r. The length of each tangled domain can then be drawn from a distribution of lengths less than the jet width. To investigate the effects of changing this distribution, we calculate best-fit $P_{\gamma\gamma}$s for 50 realizations of the tangled field (fixing $\alpha = 1$ and $r_T = r_{tr}$) in the jet of 3C454.3, for four different l_c distributions: uniform, normal (with $\langle l_c \rangle = R/2$), linearly ascending ($\propto l_c$), and linearly descending ($\propto l_c^{-1}$). Over the $N_r = 50$ realizations, the average $P_{\gamma\gamma}$ differences between a uniform l_c distribution and the others, weighted by *Fermi* expected counts,

$$\delta_{\mathrm{FW}} = \frac{1}{N_r} \sum_j \frac{\sum_i |\Delta P_{\gamma\gamma}(i,j)| \mu_i}{\sum_i \mu_i}, \qquad (A1)$$

is below 2.5% for all sources. Figure 10 shows these differences between the uniform and the normal distributions; differences between all the other distributions look similar. This means that the differences between separate realizations of the tangled field outweigh the differences between specific distributions of tangled coherence lengths. We therefore do not need to model the l_c distribution in detail and can fix it to a uniform distribution with $l_c < R$.

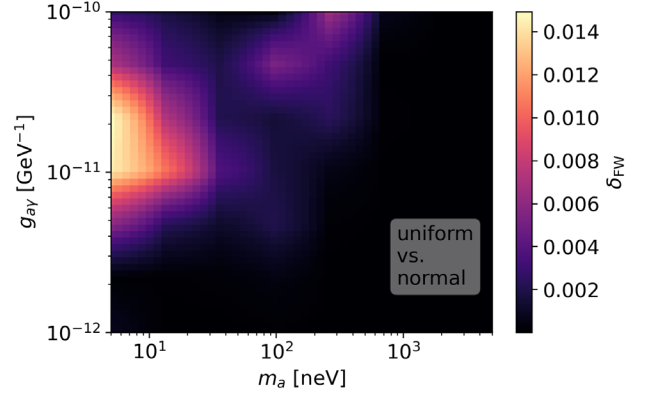The precise values of r_T and α are hard to constrain observationally, but we can test their effects on the

FIG. 10. Average best-fit $P_{\gamma\gamma}$ differences (weighted by *Fermi* counts) δ_{FW} for 3C454.3, over 50 realizations between a uniform distribution of tangled coherence lengths (l_c) and a normal distribution. Differences are <2.5%.

oscillations with the same average $P_{\gamma\gamma}$ method. When $\alpha = 1$ is fixed, the differences between $r_T = 0.3$ pc and $r_T = r_{tr}$ (again for 3C454.3) are below 1.5%. It seems that the differences in the ordered component of the field produced by varying r_T are swamped by the differences between separate realizations of the tangled component, and also, B_0 can vary in the fit to compensate for any changes. Therefore, it is reasonable to take $r_T = r_{tr}$ for each of the three sources. Figure 11 shows δ_{FW} for $\alpha = 0.5$ and $\alpha = 1$ with $r_T = r_{tr}$, when B_0 is left free in the fit, again for 3C454.3. The differences in this case can be larger ($\sim 10\%$) but only in a few isolated regions. Note that α has a larger effect than r_T because varying it can produce a larger change in the transverse field strength at the emission region. It is, of course, possible that these relatively large percentage differences in the $P_{\gamma\gamma}$s will not greatly affect the final results because they do not occur in important regions of parameter space. We perform the analysis using $\alpha = 1$ throughout. In order to test the effects of changing α on our final results, we perform a single analysis of 3C454.3 and
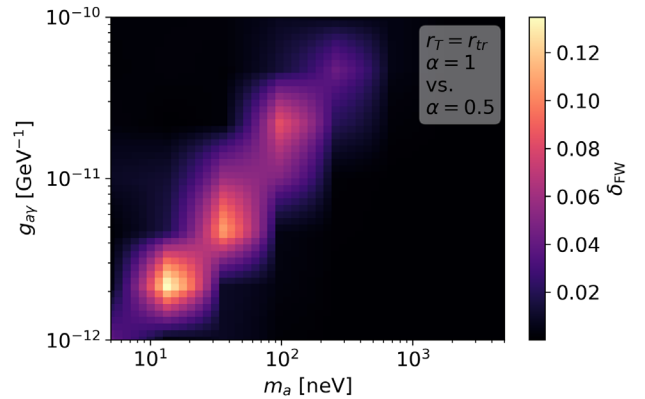

FIG. 11. Average best-fit $P_{\gamma\gamma}$ differences (weighted by *Fermi* counts) δ_{FW} for 3C454.3, over 50 realizations for different jet field parameters: $\alpha = 0.5$ and $\alpha = 1$ with $r_T = r_{tr}$ (top).
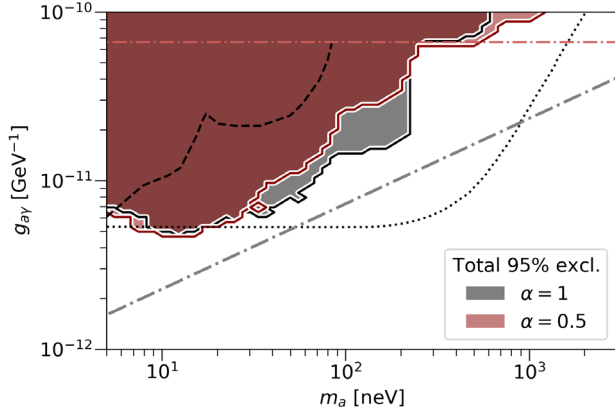
FIG. 12. Overall 95% exclusion contours for the combined analysis (black) and the same when $\alpha = 0.5$ (red). The black dotted contour shows constraints from magnetic white dwarf radio polarization [62]. Black dashed contours show previous gamma-ray constraints [20,21]. The red dot-dash contour shows the CAST experimental constraints [59]. The dark matter line is shown in grey dot-dash.

CTA 102 with $\alpha = 0.5$, calculating $\lambda(m_a, g_{a\gamma})$ in the same way as described in Sec. IV. Figure 12 shows how the total (combined) exclusions would vary in this case, using the same $\lambda_{\rm tot}^{95}$ threshold calculated with the ordinary analysis. As can be seen, the overall 95% exclusions are only slightly changed by changing α to 0.5; in some places the exclusions are slightly better, in some places they are slightly worse. This is because the regions where α can make a large difference to the $P_{\gamma\gamma}$s are generally beyond, or at the edge of, our exclusion region. This means that, particularly at lower masses ($m_a \lesssim 200$ neV), our new exclusions are robust despite the approximations made concerning the magnetic field structure.

Overall, then, the values we choose are $\alpha = 1$ and $r_T = r_{\rm tr}$ for all our sources. Figure 2 shows one example field realization for 3C454.3 with $f = 0.3$, $\alpha = 1$, and $r_T = r_{\rm tr} = 59.8$ pc.

APPENDIX B: SELF-CONSISTENCY OF FIELD AND JET MODELS

In order to check the self-consistency of our jet and background field models, we calculate steady-state and flaring SEDs for each of our sources with the AGNPY Python package,[13] and compare them to broadband observations. We follow the same method as Ref. [53]. This is not supposed to be a detailed SED modeling of our sources (indeed, the spectral parameters and magnetic field strength will be left free in the actual fits), but rather a check that our overall source models are reasonable. To calculate these SEDs, we line up spherical plasma blobs down the jet, each with a field strength (within σ_B, the errors derived from

TABLE IV. Parameters used for the blobs down the jet. The steady-state and flaring gamma-ray emission is produced from blobs at $r_{\rm vhe}$ and $r_{\rm em}$, respectively.

Parameter	$r_{\rm ss}$	$r_{\rm em}$	Rest of jet
		3C454.3	
r (pc)	59.8	0.103	
B (G)	0.013	1.9	
n_e (cm^{-3})	4.71	7.9×10^3	
E_c (MeV)	1.3×10^3	250	1.3
β	1.95	2	2
$\eta_{\rm em}$		2.8	
		CTA 102	
r (pc)	56.6	0.104	
B (G)	0.026	3.64	
n_e (cm^{-3})	2.5	3.7×10^3	
E_c (MeV)	1.3×10^3	350	1
β	1.68	1.92	2
$\eta_{\rm em}$		1.8	
		3C279	
r (pc)	47.9	0.016	
B (G)	0.0063	2.85	
n_e (cm^{-3})	5	5.25×10^4	
E_c (MeV)	1.6×10^3	380	1.3
β	1.75	2.05	2
$\eta_{\rm em}$		2.5	

Ref. [55]), electron density, and bulk Lorentz factor taken from our global jet models (see Table II). Every blob contains a population of electrons with a power-law distribution function in energy, up to a cutoff: $N_e(E) = \kappa E^{-\beta} \exp(-E/E_c)$. The synchrotron emission from all these blobs can be calculated. We can also accelerate electrons (by increasing E_c and adjusting β) within individual blobs and calculate their synchrotron and inverse-Compton emission using our field models[14] to simulate localized gamma-ray emission regions. For the steady-state emission, we accelerate electrons in the blob at $r_{\rm ss} = r_{\rm tr}$, as expected from the PC framework, where acceleration is due to a permanent large-scale feature of the jet (e.g., a standing shock, though the detailed acceleration mechanism is not modeled here or in the PC framework). For the flare emission regions, we use a blob located within the jet at $r_{\rm em}$, with a radius $R_{\rm em} = R/\eta_{\rm em}$, where R is the jet width and $\eta_{\rm em} > 1$. This roughly simulates, e.g., a reconnection or magnetoluminescence event within the highly-magnetised region of the jet, as opposed to a large-scale change in jet structure for the flares, which is disfavored because of the small flaring timescales. Table IV shows the parameters used for the

---

[14] With a ring torus model, as implemented in AGNPY, placed at the center of our tori, as opposed to our elliptical cross section model. All other field models are the same.
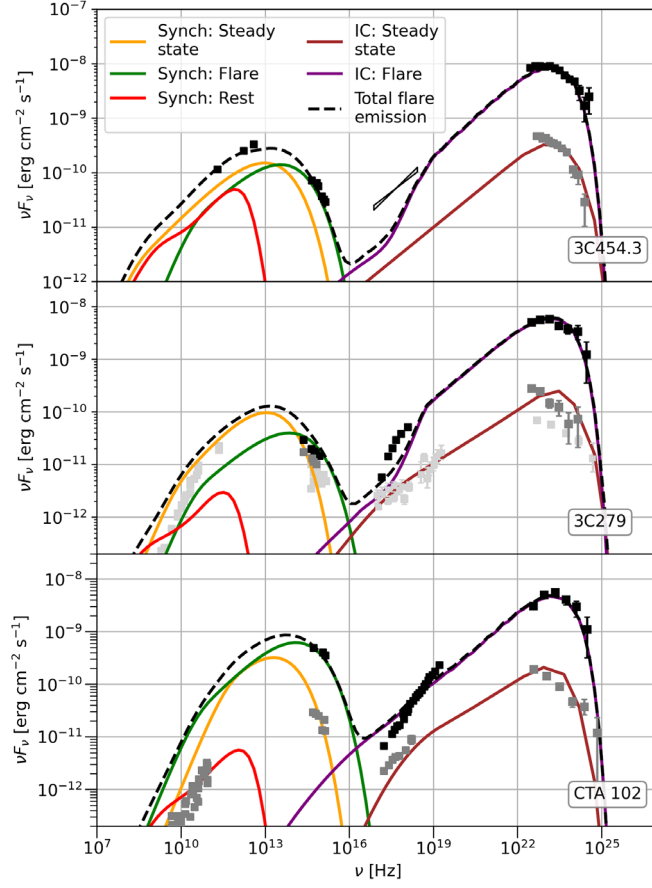
FIG. 13. Modeled SEDs for our sources during flare and steady states: 3C454.3 (top), 3C279 (middle), CTA 102 (bottom). Data for 3C454.3 taken from [68], for 3C279 from [54,57], and for CTA 102 from [69].

various blobs. Figure 13 shows our model SEDs; they are largely consistent with observations in both the flaring and steady-state cases, so we can have confidence in our overall jet and field models. This process also enables us to calculate the synchrotron fields within our jets, using the same method as Ref. [53]. Each point r will see an isotropic distribution of synchrotron photons from the surrounding blob and an anisotropic one from all the other blobs in the jet. Dispersion off the synchrotron field within the jet is always subdominant, however, and so we do not recalculate it every time B_0 changes in the fits.

APPENDIX C: SYSTEMATICS

There are also systematic uncertainties associated with the *Fermi* instrument response function, which, as pointed out in Ref. [48], could be important for ALP searches because of the spectral resolution they require. In particular, we are concerned with uncertainties associated with the energy dispersion and reconstruction. Energy dispersion is included in the analysis within FERMIPY, with different detector response matrices (DRMs) associated with each
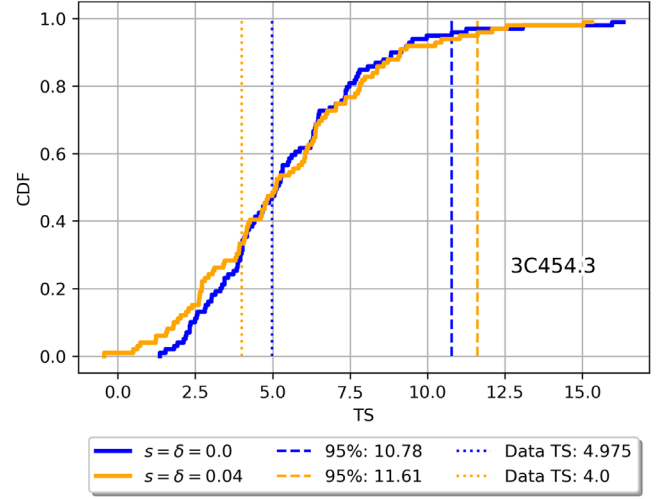

FIG. 14. 3C454.3 TS distribution with (orange) and without (blue) additional systematics ($s = \delta = 0.04$). Vertical dashed lines show the 95% thresholds; vertical dotted lines show the TS value of the data.

EDISP class (see Sec. II). There are, however, slight uncertainties in these DRMs. In particular, there could be an additional shift and smearing in the reconstructed energy and energy dispersion. These uncertainties can be included with a new expression for the expected counts [48]:

$$\mu(m_a, g_{a\gamma}, E) = \frac{1}{\mathcal{N}} \int_0^\infty dE' \exp\left(-\frac{(E - E')^2}{2(\delta E)^2}\right)$$
$$\times \mu((1 - s)E', \theta)$$
$$\times P_{\gamma\gamma}(m_a, g_{a\gamma}, \mathbf{B}, (1 - s)E'), \quad \text{(C1)}$$

where the parameters s and δ deal with an energy shift and smear, respectively, and

$$\mathcal{N} = \int_0^\infty dE' \exp\left(-\frac{(E - E')^2}{2(\delta E)^2}\right). \quad \text{(C2)}$$

Unfortunately, having to perform this integral at every fine energy that we calculate $P_{\gamma\gamma}$ at greatly increases the computation time, and so it is not feasible to perform the whole analysis in this way—especially as we leave B_0 free in the fitting.

These errors are expected to be $\lesssim$ a few percent for *Fermi*.[15] Therefore, to test the possible effects of s and δ on our final results, we perform a single analysis of 3C454.3 using $s = \delta = 0.04$.

Changes in the systematics would also likely affect the λ_{95} thresholds, so the comparison used in Fig. 12, with the same thresholds as the regular analysis, is less useful in this

---

[15]See https://fermi.gsfc.nasa.gov/ssc/data/analysis/LAT_caveats.html as accessed on October 5, 2022.

case. Nevertheless, we can compare the TS distributions for the two cases (with and without additional systematics) because the smooth no-ALP spectrum should be unaffected. Figure 14 shows TS distributions, for the $s = \delta = 0$ (blue) and $s = \delta = 0.04$ (orange) cases, for 100 simulated data sets that do not include an injected ALP signal. The vertical dashed lines show the 95% thresholds, and the vertical dotted lines show the TS values of the data. As can be seen, the TS distributions are very similar in the two cases. An extra shift and (particularly) a smear in the energy reconstruction reduces our ability to distinguish the ALP

and no-ALP cases, both slightly reducing the TS value of the data and slightly increasing the 95% threshold. Therefore, we would expect the additional systematics to slightly shrink our exclusion regions (in [21], they show that including similar systematic errors would reduce their exclusion region by ∼6%). This is to be expected, as the Gaussian (δ) term effectively smooths out the oscillations. Nevertheless, even in this case, our observations would likely still constrain ALP parameter space previously unprobed by gamma-ray searches—though more computing power would be required to perform the full analysis.

[1] R. D. Peccei and Helen R. Quinn, *CP* Conservation in the Presence of Pseudoparticles, Phys. Rev. Lett. **38,** 1440 (1977).

[2] Steven Weinberg, A New Light Boson?, Phys. Rev. Lett. **40,** 223 (1978).

[3] F. Wilczek, Problem of Strong *p* and *t* Invariance in the Presence of Instantons. Phys. Rev. Lett. **40,** 279 (1978).

[4] R. D. Peccei, The strong *CP* problem and axions, Lect. Notes Phys. **741,** 3 (2008).

[5] K. Nakamura, Review of particle physics, J. Phys. G **37,** 075021 (2010).

[6] Francesca Chadha-Day, John Ellis, and David J. E. Marsh, Axion dark matter: What is it and why now?, Sci. Adv. **8,** eabj3618 (2022).

[7] Yannis K. Semertzidis and SungWoo Youn, Axion dark matter: How to see it?, Sci. Adv. **8,** eabm9928 (2022).

[8] Neil Turok, Almost-Goldstone Bosons from Extra-Dimensional Gauge Theories, Phys. Rev. Lett. **76,** 1015 (1996).

[9] Joerg Jaeckel and Andreas Ringwald, The low-energy frontier of particle physics, Annu. Rev. Nucl. Part. Sci. **60,** 405 (2010).

[10] Andreas Ringwald, Searching for axions and ALPs from string theory, J. Phys. Conf. Ser. **485,** 012013 (2014).

[11] Igor G. Irastorza and Javier Redondo, New experimental approaches in the search for axion-like particles, Prog. Part. Nucl. Phys. **102,** 89 (2018).

[12] John Preskill, Mark B. Wise, and Frank Wilczek, Cosmology of the invisible axion, Phys. Lett. B **120,** 127 (1983).

[13] L. F. Abbott and P. Sikivie, A cosmological bound on the invisible axion, Phys. Lett. B **120,** 133 (1983).

[14] Michael Dine and Willy Fischler, The not-so-harmless axion, Phys. Lett. B **120,** 137 (1983).

[15] Paola Arias, Davide Cadamuro, Mark Goodsell, Joerg Jaeckel, Javier Redondo, and Andreas Ringwald, WISPy cold dark matter, J. Cosmol. Astropart. Phys. 06 (2012) 013.

[16] Peter W. Graham, Igor G. Irastorza, Steven K. Lamoreaux, Axel Lindner, and Karl A. van Bibber, Experimental searches for the axion and axion-like particles, Annu. Rev. Nucl. Part. Sci. **65,** 485 (2015).

[17] J. Isern, E. García-Berro, S. Torres, and S. Catalán, Axions and the cooling of white dwarf stars, Astrophys. J. Lett. **682,** L109 (2008).

[18] Alexandre Payez, Carmelo Evoli, Tobias Fischer, Maurizio Giannotti, Alessandro Mirizzi, and Andreas Ringwald, Revisiting the SN1987A gamma-ray limit on ultralight axion-like particles, J. Cosmol. Astropart. Phys. 02 (2015) 006.

[19] Dong Lai and Jeremy Heyl, Probing axions with radiation from magnetic stars, Phys. Rev. D **74,** 123003 (2006).

[20] A. Abramowski *et al.*, Constraints on axionlike particles with H.E.S.S. from the irregularity of the PKS 2155-304 energy spectrum, Phys. Rev. D **88,** 102003 (2013).

[21] M. Ajello *et al.*, Search for Spectral Irregularities due to Photon–Axionlike-Particle Oscillations with the Fermi Large Area Telescope, Phys. Rev. Lett. **116,** 161101 (2016).

[22] Christopher S. Reynolds, M. C. David Marsh, Helen R. Russell, Andrew C. Fabian, Robyn Smith, Francesco Tombesi, and Sylvain Veilleux, Astrophysical limits on very light axion-like particles from chandra grating spectroscopy of NGC 1275, Astrophys. J. **890,** 59 (2020).

[23] R. Buehler, G. Gallardo, G. Maier, A. Domínguez, M. López, and M. Meyer, Search for the imprint of axion-like particles in the highest-energy photons of hard γ-ray blazars, J. Cosmol. Astropart. Phys. 09 (2020) 027.

[24] James H. Matthews, Christopher S. Reynolds, M. C. David Marsh, Júlia Sisk-Reynés, and Payton E. Rodman, How do magnetic field models affect astrophysical limits on light axion-like particles? An X-ray case study with NGC 1275, Astrophys. J. **930,** 90 (2022).

[25] Júlia Sisk-Reynés, James H. Matthews, Christopher S. Reynolds, Helen R. Russell, Robyn N. Smith, and M. C. David Marsh, New constraints on light axion-like particles using Chandra transmission grating spectroscopy of the powerful cluster-hosted quasar H1821 + 643, Mon. Not. R. Astron. Soc. **510,** 1264 (2022).

[26] S. Abdollahi *et al.*, Fermi large area telescope fourth source catalog, Astrophys. J. Suppl. Ser. **247,** 33 (2020).

[27] Roger Blandford, David Meier, and Anthony Readhead, Relativistic jets from active galactic nuclei, Annu. Rev. Astron. Astrophys. **57,** 467 (2019).

[28] M. Petropoulou and A. Mastichiadis, On proton synchrotron blazar models: The case of quasar 3C 279, Mon. Not. R. Astron. Soc. **426,** 462 (2012).

[29] A. Mücke, R. J. Protheroe, R. Engel, J. P. Rachen, and T. Stanev, BL Lac objects in the synchrotron proton blazar model, Astropart. Phys. **18**, 593 (2003).

[30] A. A. Zdziarski and M. Bottcher, Hadronic models of blazars require a change of the accretion paradigm, Mon. Not. R. Astron. Soc. **450**, L21 (2015).

[31] Enrico Fermi, On the origin of the cosmic radiation, Phys. Rev. **75**, 1169 (1949).

[32] Alan P. Marscher, Turbulent, extreme multi-zone model for simulating flux and polarization variability in blazars, Astrophys. J. **780**, 87 (2014).

[33] Alessandro de Angelis, Giorgio Galanti, and Marco Roncadelli, Relevance of axionlike particles for very-high-energy astrophysics, Phys. Rev. D **84**, 105030 (2011).

[34] Kathrin A. Hochmuth and Günter Sigl, Effects of axion-photon mixing on gamma-ray spectra from magnetized astrophysical sources, Phys. Rev. D **76**, 123011 (2007).

[35] Nicola Bassan and Marco Roncadelli, Photon-axion conversion in active galactic nuclei?, arXiv:0905.3752.

[36] M. A. Sanchez-Conde, D. Paneque, E. Bloom, F. Prada, and A. Dominguez, Hints of the existence of axion-like-particles from the gamma-ray spectra of cosmological sources, Phys. Rev. D **79**, 123511 (2009).

[37] M. Fairbairn, T. Rashba, and S. Troitsky, Photon-axion mixing and ultra-high energy cosmic rays from BL Lac type objects: Shining light through the Universe, Phys. Rev. D **84**, 125019 (2011).

[38] Jonathan Harris and Paula M. Chadwick, Photon-axion mixing within the jets of active galactic nuclei and prospects for detection, J. Cosmol. Astropart. Phys. 10 (2014) 018.

[39] Fabrizio Tavecchio, Marco Roncadelli, and Giorgio Galanti, Photons to axion-like particles conversion in active galactic nuclei, Phys. Lett. B **744**, 375 (2015).

[40] James Davies, Manuel Meyer, and Garret Cotter, Relevance of jet magnetic field structure for blazar axionlike particle searches, Phys. Rev. D **103**, 023008 (2021).

[41] W. B. Atwood et al., The large area telescope on the Fermi gamma-ray space telescope mission, Astrophys. J. **697**, 1071 (2009).

[42] Manuel Meyer, Jeffrey D. Scargle, and Roger D. Blandford, Characterizing the gamma-ray variability of the brightest flat spectrum radio quasars observed with the Fermi LAT, Astrophys. J. **877**, 39 (2019).

[43] M. Wood, R. Caputo, E. Charles, M. Di Mauro, J. Magill, J. S. Perkins (Fermi-LAT Collaboration), FERMIPY: An open-source Python package for analysis of Fermi-LAT data, Proc. Sci. ICRC2017 (**2017**) 824.

[44] W. Atwood, A. Albert, L. Baldini, M. Tinivella, J. Bregeon, M. Pesce-Rollins, C. Sgrò, P. Bruel, E. Charles, A. Drlica-Wagner, A. Franckowiak, T. Jogler, L. Rochester, T. Usher, M. Wood, J. Cohen-Tanugi, and S. Zimmer, Pass 8: Toward the full realization of the Fermi-LAT scientific potential, arXiv:1303.3514.

[45] Georg Raffelt and Leo Stodolsky, Mixing of the photon with low-mass particles, Phys. Rev. D **37**, 1237 (1988).

[46] Alexandra Dobrynina, Alexander Kartavtsev, and Georg Raffelt, Photon-photon dispersion of Tev gamma rays and its role for photon-ALP conversion, Phys. Rev. D **91**, 083003 (2015).

[47] Ronnie Jansson and Glennys R. Farrar, A new model of the galactic magnetic field, Astrophys. J. **757**, 14 (2012).

[48] H. Abdalla et al., Sensitivity of the Cherenkov telescope array for probing cosmology and fundamental physics with gamma-ray propagation, J. Cosmol. Astropart. Phys. 02 (2021) 048.

[49] A. Dominguez et al., Extragalactic background light inferred from AEGIS galaxy SED-type fractions, Proc. Sci. TEXAS2010 (**2010**) 228.

[50] Manuel Meyer, James Davies, and Julian Kuhlmann, gammaALPs, v0.2.0, 10.5281/zenodo.4973513 (2021).

[51] James Davies, Manuel Meyer, and Garret Cotter, Data files for Constraints on axionlike particles from a combined analysis of three flaring Fermi flat-spectrum radio quasars, Zenodo, 10.5281/zenodo.7573123 (2023).

[52] Manuel Meyer, James Davies, and Julian Kuhlmann, gammaALPs: An open-source Python package for computing photon-axion-like-particle oscillations in astrophysical environments, Proc. Sci. ICRC2021 (**2021**) 557.

[53] James Davies, Manuel Meyer, and Garret Cotter, Relevance of photon-photon dispersion within the jet for blazar axionlike particle searches, Phys. Rev. D **105**, 023017 (2022).

[54] William J. Potter and Garret Cotter, New constraints on the structure and dynamics of black hole jets, Mon. Not. R. Astron. Soc. **453**, 4070 (2015).

[55] Mohammad Zamaninasab, E Clausen-Brown, T Savolainen, and A Tchekhovskoy, Dynamically important magnetic fields near accreting supermassive black holes, Nature (London) **510**, 126 (2014).

[56] Justin D. Finke, External compton scattering in blazar jets and the location of the gamma-ray emitting region, Astrophys. J. **830**, 94 (2016).

[57] H. Abdalla et al. (H. E. S. S. Collaboration), Constraints on the emission region of 3C 279 during strong flares in 2014 and 2015 through VHE γ-ray observations with H.E.S.S, Astron. Astrophys. **627**, A159 (2019).

[58] G. Calderone, T. Sbarrato, and G. Ghisellini, The wide-field infrared survey explorer view of the disc-torus connection in z ~0.6 active galactic nuclei, Mon. Not. R. Astron. Soc. **425**, L41 (2012).

[59] CAST Collaboration, New CAST limit on the axion-photon interaction, Nat. Phys. **13**, 584 (2017).

[60] Wolfgang A. Rolke, Angel M. López, and Jan Conrad, Limits and confidence intervals in the presence of nuisance parameters, Nucl. Instrum. Methods Phys. Res., Sect. A **551**, 493 (2005).

[61] M. Ackermann et al. (Fermi-LAT Collaboration), Searching for Dark Matter Annihilation from Milky Way Dwarf Spheroidal Galaxies with Six Years of Fermi Large Area Telescope Data, Phys. Rev. Lett. **115**, 231301 (2015).

[62] Christopher Dessert, David Dunsky, and Benjamin R. Safdi, Upper limit on the axion-photon coupling from magnetic white dwarf polarization, Phys. Rev. D **105**, 103034 (2022).

[63] M. Diaz Ortiz, J. Gleason, H. Grote, A. Hallal, M. T. Hartman, H. Hollis, K. S. Isleif, A. James, K. Karan, T. Kozlowski, A. Lindner, G. Messineo, G. Mueller, J. H. Poeld, R. C. G. Smith, A. D. Spector, D. B. Tanner,

L. W. Wei, and B. Willke, Design of the ALPS II optical system, Phys. Dark Universe **35**, 100968 (2022).

[64] E. Armengaud *et al.*, Conceptual design of the international axion observatory (IAXO), J. Instrum. **9**, 01 (2014).

[65] M. Zamaninasab, T. Savolainen, E. Clausen-Brown, T. Hovatta, M. L. Lister, T. P. Krichbaum, Y. Y. Kovalev, and A. B. Pushkarev, Evidence for a large-scale helical magnetic field in the quasar 3C 454.3, Mon. Not. R. Astron. Soc. **436**, 3341 (2013).

[66] Xiaofeng Li, P. Mohan, T. An, Xiaoyu Hong, Xiaopeng Cheng, Jun Yang, Yingkang Zhang, Zhongli Zhang, and Wei Zhao, Imaging and variability studies of CTA 102

during the 2016 January γ-ray flare, Astrophys. J. **854**, 17 (2018).

[67] D. C. Homan, M. L. Lister, H. D. Aller, M. F. Aller, and J. F. C. Wardle, Full polarization spectra of 3C 279, Astrophys. J. **696**, 328 (2009).

[68] Matteo Cerruti, Charles D. Dermer, Benoît Lott, Catherine Boisson, and Andreas Zech, Gamma-ray blazars near equipartition and the origin of the GeV spectral break in 3C 454.3, Astrophys. J. Lett. **771**, L4 (2013).

[69] S. Gasparyan, N. Sahakyan, V. Baghmanyan, and D. Zargaryan, On the multiwavelength emission from CTA 102, Astrophys. J. **863**, 114 (2018).