

Constraints on axionlike polarization oscillations in the cosmic microwave background with POLARBEAR

Shunsuke Adachi^{1,2}, Tylor Adkins³, Kam Arnold⁴, Carlo Baccigalupi^{5,6,7}, Darcy Barron⁸, Kolen Cheung^{3,9,10}, Yuji Chinone^{11,12,*}, Kevin T. Crowley⁴, Josquin Errard¹³, Giulio Fabbian^{14,15}, Chang Feng¹⁶, Raphael Flauger⁴, Takuro Fujino¹⁷, Daniel Green⁴, Masaya Hasegawa^{18,11,19}, Masashi Hazumi^{11,18,20,12}, Daisuke Kaneko¹¹, Nobuhiko Katayama¹², Brian Keating⁴, Akito Kusaka^{21,22,12,23}, Adrian T. Lee^{3,21}, Yuto Minami²⁴, Haruki Nishino²³, Christian L. Reichardt²⁵, Praween Siritanasak²⁶, Jacob Spisak^{4,†}, Osamu Tajima², Satoru Takakura²⁷, Sayuri Takatori²⁸, Grant Paul Teply⁴, and Kyohei Yamada²²

(POLARBEAR Collaboration)

¹*Hakubi Center for Advanced Research, Kyoto University, Kyoto 606-8501, Japan*

²*Department of Physics, Faculty of Science, Kyoto University, Kyoto 606-8502, Japan*

³*Department of Physics, University of California, Berkeley, Berkeley, California 94720, USA*

⁴*Department of Physics, University of California, San Diego, La Jolla, California 92093, USA*

⁵*Astrophysics and Cosmology, The International School for Advanced Studies (SISSA), Trieste 34136, Italy*

⁶*Astrophysics and Cosmology, Institute for Fundamental Physics of the Universe (IFPU), Trieste 34151, Italy*

⁷*Astrophysics and Cosmology, The National Institute for Nuclear Physics (INFN), Trieste 34127, Italy*

⁸*Department of Physics and Astronomy, University of New Mexico, Albuquerque, New Mexico 87131, USA*

⁹*Space Sciences Laboratory, University of California, Berkeley, Berkeley, California 94720, USA*

¹⁰*Computational Cosmology Center, Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA*

¹¹*QUP (WPI), KEK, Tsukuba, Ibaraki 305-0801, Japan*

¹²*Kavli Institute for the Physics and Mathematics of the Universe (WPI), UTIAS, The University of Tokyo, Kashiwa, Chiba 277-8583, Japan*

¹³*Université de Paris Cité, CNRS, Astroparticule et Cosmologie, F-75013 Paris, France*

¹⁴*Center for Computational Astrophysics, Flatiron Institute, New York, New York 10003, USA*

¹⁵*School of Physics and Astronomy, Cardiff University, Cardiff CF24 3AA, United Kingdom*

¹⁶*Department of Astronomy, School of Physical Sciences, University of Science and Technology of China, Hefei 230026, China*

¹⁷*Graduate School of Engineering Science, Yokohama National University, Yokohama 240-8501, Japan*

¹⁸*Institute of Particle and Nuclear Studies (IPNS), High Energy Accelerator Research Organization (KEK), Tsukuba 305-0801, Japan*

¹⁹*SOKENDAI, SOKENDAI, Tsukuba 305-0801, Japan*

²⁰*Institute of Space and Astronautical Science (ISAS), Japan Aerospace Exploration Agency (JAXA), Sagami 252-5210, Japan*

²¹*Physics Division, Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA*

²²*Department of Physics, The University of Tokyo, Tokyo 113-0033, Japan*

²³*Research Center for the Early Universe, Graduate School of Science, The University of Tokyo, Tokyo 113-0033, Japan*

²⁴*Research Center for Nuclear Physics, Osaka University, Ibaraki 567-0047, Japan*

²⁵*School of Physics, The University of Melbourne, Parkville 3010, Australia*

²⁶*National Astronomical Research Institute of Thailand, Chiangmai 50180, Thailand*

²⁷*Department of Astrophysical and Planetary Sciences, University of Colorado, Boulder, Colorado 80309, USA*

²⁸*Research Institute for Interdisciplinary Science (RIIS), Okayama University, Okayama 700-8530, Japan*



(Received 16 March 2023; accepted 26 June 2023; published 18 August 2023)

*Corresponding author: chinoney@gmail.com

†Corresponding author: jspisak@ucsd.edu

Very light pseudoscalar fields, often referred to as axions, are compelling dark matter candidates and can potentially be detected through their coupling to the electromagnetic field. Recently a novel detection technique using the cosmic microwave background (CMB) was proposed, which relies on the fact that the axion field oscillates at a frequency equal to its mass in appropriate units, leading to a time-dependent birefringence. For appropriate oscillation periods this allows the axion field at the telescope to be detected via the induced sinusoidal oscillation of the CMB linear polarization. We search for this effect in two years of POLARBEAR data. We do not detect a signal, and place a median 95% upper limit of 0.65° on the sinusoid amplitude for oscillation frequencies between 0.02 days^{-1} and 0.45 days^{-1} , which corresponds to axion masses between $9.6 \times 10^{-22} \text{ eV}$ and $2.2 \times 10^{-20} \text{ eV}$. Under the assumptions that (1) the axion constitutes all the dark matter and (2) the axion field amplitude is a Rayleigh-distributed stochastic variable, this translates to a limit on the axion-photon coupling $g_{\phi\gamma} < 2.4 \times 10^{-11} \text{ GeV}^{-1} \times (m_\phi/10^{-21} \text{ eV})$.

DOI: [10.1103/PhysRevD.108.043017](https://doi.org/10.1103/PhysRevD.108.043017)

I. INTRODUCTION

The nature of dark matter, particularly its nongravitational interactions, remains one of the biggest open questions in cosmology and particle physics. One possibility that has recently received significant attention is low-mass bosonic dark matter [1–3]. The canonical example is the original QCD axion, a pseudo-Nambu Goldstone boson associated with the spontaneous breaking of a U(1) symmetry proposed to solve the strong CP problem. [4–7]. More generally, a broader class of pseudoscalar fields with small masses and couplings to the standard model have been considered. These are often called axionlike-particles and are not necessarily solutions to the strong CP problem; nevertheless, we will refer to them as axions in this work for brevity. String theory generically predicts the existence of many such axions populating a wide range of masses and couplings, which is sometimes called the axiverse [8,9]. Axions with astrophysically large de Broglie wavelengths ($\lambda \sim 1 \text{ kpc}$ for axion mass $m_\phi \sim 10^{-22} \text{ eV}$) are a particularly intriguing dark matter candidate because of their ability to act as fuzzy dark matter, which can potentially resolve conflicts between small-scale predictions of cold dark matter models and observations [10,11].

One way to detect axions is via their interaction with electromagnetism. Laboratory experiments such as ADMX [12] and ABRACADABRA [13], for example, exploit the coupling between axions and magnetic fields to set limits on the axion dark matter in the QCD mass range. Reference [14] provides an overview of experimental approaches. It is well known that the coupling between the electromagnetic field and a pseudoscalar field generates an effective birefringence, rotating linearly polarized light [15,16]. The axion is well modeled as a classical field, which, when the axion mass is less than the Hubble rate, oscillates at a frequency equal to its mass. Fedderke *et al.* [17], hereafter F19, pointed out two novel effects in the cosmic microwave background (CMB) caused by this oscillation and birefringence. The first is the suppression

of the overall CMB polarization signal due to averaging over many axion oscillation periods during recombination, which was constrained in F19. The second is a coherent, all-sky oscillation of the CMB’s linear polarization due to the oscillation of the axion field *at the telescope*. This effect can be constrained when the oscillation period is appropriate for CMB experiments, e.g. hours to years, corresponding to masses in the 10^{-19} eV – 10^{-22} eV range. This also happens to be the mass range in which the axion can act as fuzzy dark matter [10]. In this work, we search for this effect using data from the POLARBEAR experiment.

This signal has been constrained by other CMB experiments: BICEP/Keck, in [18,19] (the latter hereafter BK22), and the South Pole Telescope (SPT) in [20] (hereafter SPT22). Our analysis is similar to these analyses, with two primary differences. The first is that we estimate the CMB polarization angle using C_ℓ^{EB} power spectra rather than in Q/U pixel space. This allows us to use the theoretical C_ℓ^{EE} power spectrum from precisely measured cosmological parameters as a polarization template rather than coadded Q/U maps, and also facilitates easier systematic error checks. The second is that we model the amplitude of the axion field at the telescope as a stochastic, Rayleigh-distributed variable rather than assuming a fixed value corresponding to the mean Milky-Way halo density [21]. The need for this approach was pointed out in [22] and weakens the median constraint on the axion-photon coupling constant by a factor of 2.2 in our analysis.

The birefringence signature generated by axion dark matter can be constrained by other astrophysical polarized sources, and many other authors have used this effect to place constraints in a similar mass range using pulsars, active galactic nuclei, protoplanetary disks, Sagittarius A*, and black hole superradiance [23–30]. The CMB, however, has several attractive features that make it ideal for this type of analysis. The signal is entirely due to the axion field at the telescope, so we do not require any modeling of the field at the source (during the release of the CMB). The template polarization signal from the CMB has minimal time-dependent contamination and is extremely well

measured across many experiments. Finally, POLARBEAR and other CMB instruments have a long history of precision CMB polarization measurements with well-understood noise properties. These factors all serve to mitigate the systematic error in this analysis.

In addition to birefringence-based measurements, bounds from cosmological structure like the Lyman- α forest and Milky Way satellites have placed constraints on the minimum allowed axion mass [31–33] if it is fuzzy dark matter. A constraint on the minimum allowed mass has also been derived from the impact of dynamical heating on the velocity dispersions of ultra-faint dwarf galaxies [34]. These bounds strongly constrain the allowed parameter space for fuzzy dark matter, but are subject to different systematic and modeling uncertainties than birefringence analyses. Upper bounds on the axion-photon coupling which are constant in the mass range we consider have been derived from axion-photon conversion in the Sun, supernova 1987A, quasar H1821 + 643 and in the intra-cluster medium [35–38].

To perform this analysis we use the first two seasons of data from the POLARBEAR experiment, which measured the polarization of the CMB during 2012–2016. During the first two seasons of observations, POLARBEAR observed three small sky patches in order to measure gravitationally lensed B -modes. These results were presented in [39] (hereafter PB14) and [40] (hereafter PB17). The telescope has an angular resolution of 3.5 arcmin and reported measurements of the B -mode power spectrum up to multipoles of $\ell = 2100$ from the first two seasons, including the angular scales at which the CMB polarization signal is the strongest. POLARBEAR measurements have previously been used to constrain anisotropic birefringence and primordial magnetic fields [41]: this work is entirely separate, however, because we are searching for *time-dependent* oscillations of the isotropic birefringence angle, whereas [41] considered time-independent birefringence.

The rest of the paper is organized as follows. In Sec. II, we describe the POLARBEAR instrument and observations. In Sec. III, we detail the analysis procedure used to generate CMB polarization angles and search for an axion signal. In Sec. IV and V, we describe the null tests and systematics estimates used to validate the dataset. The results are presented in Sec. VI, and the conclusion in Sec. VII.

II. FIRST AND SECOND SEASON OBSERVATIONS OF THE POLARBEAR INSTRUMENT

The POLARBEAR experiment consisted of a cryogenic receiver attached to a two-mirror reflective telescope, the 2.5 m Huan Tran Telescope. It was located at the James Ax Observatory in the Atacama Desert in Chile at an elevation of 5,190 m. The receiver contained 1,274 transition-edge sensors arranged into 637 polarization-sensitive pixels situated in 7 detector wafers on the focal plane, which was cooled to 0.3 K, and observed at a single frequency

centered at 150 GHz. We will report data from the first and second season observations, which occurred from May 2012 to June 2013 and October 2013 to April 2014, respectively. More details about the POLARBEAR receiver and telescope can be found in [42] and [43].

The POLARBEAR observing strategy during the first and second seasons is described in detail in PB14 and PB17, and will be summarized here. Three separate sky patches were observed, each with an effective sky area of 7–9 square degrees. In this work, an “observation” will refer to a single, continuous measurement of one patch which lasts until the patch is no longer visible, typically 4–8 hours. Each observation consists of many 15-minute constant elevation scans (CESs). During one CES, the telescope scans back and forth in azimuth repeatedly at a constant elevation, and as the sky rotates, the entire patch is observed. The telescope then changes elevation and repeats the process. A typical observing day involves sequentially observing all three patches. This observing strategy allows us to probe oscillation periods longer than about 2 days.

During the first two seasons the telescope observed with a cryogenic half-wave plate (HWP) located on the sky-side of the lenses. During the first season it was periodically rotated between observations in order to mitigate systematic errors, and during the second season it remained fixed. In PB17, it was noted that this generated some uncertainty in the absolute polarization angle of the instrument during the first season and between the two seasons. This source of systematic error limits the maximum oscillation period we can assess to 50 days, as discussed in Sec. V.

III. ANALYSIS METHOD

A. Expected signal

Following F19, the axion-photon coupling in the Lagrangian can be written as

$$\mathcal{L} = -\frac{1}{4}g_{\phi\gamma}\phi F_{\mu\nu}\tilde{F}^{\mu\nu}, \quad (1)$$

where ϕ is the axion field, $g_{\phi\gamma}$ is the axion-photon coupling, $F_{\mu\nu}$ is the electromagnetic field tensor and $\tilde{F}^{\mu\nu}$ is its dual. Assuming that the amplitude is small enough so the potential is well-approximated by $V(\phi) = m_\phi^2\phi^2/2$, the axion field at the telescope is well described by $\phi(t) = \phi_0 \sin(m_\phi t + \theta)$. We treat the amplitude ϕ_0 and phase θ as constants because the duration of the experiment is much less than the axion coherence time, as discussed in Sec. VI C. The polarization angle of the CMB due to the axion field at the telescope is then (F19)

$$\beta_{\text{CMB}}(t) = \frac{g_{\phi\gamma}\phi_0}{2} \sin(m_\phi t + \theta). \quad (2)$$

This can be conveniently parametrized as

$$\beta_{\text{CMB}}(t) = A \sin(2\pi ft + \theta). \quad (3)$$

This is a sinusoid with unknown amplitude A , frequency f , and phase θ . The basic unit of time in the POLARBEAR survey schedule is one observation of a single patch, which can last up to 8 hours. For each observation, therefore, we construct CMB maps and estimate a single angle, generating several hundred angles over the course of the two years of data. We will then form a likelihood to search for the presence of a sinusoidal signal in this data.

B. Angle estimation procedure

There are multiple ways to estimate a rotation angle from the observation maps. The most direct way is to search for the rotation in the Stokes Q and U parameter maps by comparing them to a template of Q and U maps containing only the unrotated CMB. This is the method employed in BK22 and SPT22. Another way is to transform the Q and U maps into E -mode and B -mode maps, then use a single-observation C_{ℓ}^{EB} spectrum to estimate a rotation angle. While the latter method has the disadvantage of requiring more computationally-intensive steps, we choose to implement it for the following reasons. The first is that we can use the extremely well-determined theoretical C_{ℓ}^{EE} from other experiments as our polarization template, rather than a Q/U template map created by coadding our observations. While the C_{ℓ}^{EE} template is affected by sample variance, this effect is not important for the ℓ range we use. The second advantage is that the power spectra approach is very similar to the approach taken in the construction of the full, coadded spectra in PB17. This lets us reuse most elements of this validated pipeline, including many of the same systematics estimates.

The method of estimating a time-independent rotation angle using the C_{ℓ}^{EB} power spectrum is well established [44], and has been used in many analyses, including PB17, to correct for an overall telescope miscalibration angle. The method we will employ to search for a *time-varying* rotation angle is similar, except that we must construct a spectrum for every observation. This will be used to estimate a single rotation angle for each observation, constant over the duration of the observation. We can then search for a time-varying signal in the timestream of these angles. For a single observation with CMB rotation angle α , in the absence of noise and foreground contamination (which are addressed in Sec. III E and Sec. V C respectively), the observed Fourier transformed E/B -mode coefficients are

$$E_{\ell m}^{\text{obs}} = \cos(2\alpha)E_{\ell m}^{\text{CMB}} - \sin(2\alpha)B_{\ell m}^{\text{CMB}} \quad (4)$$

$$B_{\ell m}^{\text{obs}} = \sin(2\alpha)E_{\ell m}^{\text{CMB}} + \cos(2\alpha)B_{\ell m}^{\text{CMB}}. \quad (5)$$

This analysis is sensitive to signals of amplitude $\approx 1^\circ$, so we use the small angle approximation

$$E_{\ell m}^{\text{obs}} = E_{\ell m}^{\text{CMB}} - 2\alpha B_{\ell m}^{\text{CMB}} \quad (6)$$

$$B_{\ell m}^{\text{obs}} = 2\alpha E_{\ell m}^{\text{CMB}} + B_{\ell m}^{\text{CMB}}. \quad (7)$$

Since the CMB is E -mode dominated, to leading order all of the rotation information is contained in the $2\alpha E_{\ell m}^{\text{CMB}}$ term in Eq. (7). In order to recover an angle from a single observation, then, we can construct a power spectrum using one B -mode map correlated with the full, coadded E -mode maps. The spectra for observation j , rotated by angle α_j , is

$$\begin{aligned} C_{\ell,j}^{EB,\text{obs}} &= \frac{1}{N} \frac{1}{(2\ell+1)} \sum_i^N \sum_m E_{\ell m,i}^{\text{obs}} (B_{\ell m,j}^{\text{obs}})^* \\ &= 2\alpha_j C_{\ell}^{EE,\text{CMB}}. \end{aligned} \quad (8)$$

The E -mode maps have been coadded over all observations $i = 1, \dots, N$. The intrinsic $C_{\ell}^{EB,\text{CMB}}$ has been set to zero, and we have neglected $C_{\ell}^{BB,\text{CMB}}$ because it is $\ll C_{\ell}^{EE,\text{CMB}}$.

C. Mapmaking

The mapmaking procedure is nearly identical to the one used in ‘‘Pipeline A’’ of PB17, and will be briefly reviewed here. To make maps of the polarized sky, the raw time-ordered data (TOD) undergo a series of quality cuts, are converted to CMB temperature units, and then differenced to form polarization timestreams. The timestreams are filtered to remove high and low frequency noise as well as scan-synchronous signals then combined with the detector pointing data to make maps. The maps are apodized with noise-weighted masks, which cover point sources as well. This procedure yields a single Q and U map for each observation. Finally, Q and U maps are transformed to E and B maps using the pure B -mode transform [45]. Due to the filtering, they are biased estimates of the true sky signal.

The chief difference between the maps used in this analysis and those in PB17 is the apodization. PB17 uses a separate mask for each observation, whereas this analysis uses one apodization mask common to every observation within each patch. This is done in order to simplify the calculation of the mode-mixing matrix, discussed in Sec. III D. However, this masking procedure does sub-optimally weight each pixel within a given observation, leading to an increase in noise. In some observations the increase in noise from edge pixels is excessive, and they are removed from the analysis if the map noise increase exceeds a certain threshold. The same cut is applied to observations that have excessive noise on one split in a null test as well, as described in Sec. IV A. The overall impact of this apodization procedure is an effective noise increase of about 15%. Improving the apodization procedure to lessen this noise hit is an area of improvement for future analyses.




FIG. 1. Left: the 515 estimated per-observation rotation angles $\hat{\alpha}_j$ used in this analysis. Right: a histogram of the angles normalized by their error, $\hat{\alpha}_j/\sigma_j$. The angles are Gaussian distributed.

D. Maps to spectra

The procedure to construct single-observation C_ℓ^{EB} spectra as outlined above is identical to PB17, with the modification that only one B -mode map is used per observation. The coadded E -mode maps and a single B -mode map are used to construct pseudospectra (\tilde{C}_ℓ), which are biased estimates of the true spectra:

$$\tilde{C}_{\ell,j}^{EB} = \frac{1}{N_\ell} \sum_{k \in \text{bin}_\ell} \tilde{m}_{jk}^{B*} \left(\frac{1}{\sum_{i \neq j} w_i} \sum_{i \neq j} w_i \tilde{m}_{ik}^E \right). \quad (9)$$

The \tilde{m}_{jk} denotes a Fourier-transformed map with 2D wave vector k and observation j . The flat-sky approximation is used due to the size of the POLARBEAR patches, and the spectra are combined into ℓ bins of width $\Delta\ell = 40$. N_ℓ is the number of wave vectors in bin ℓ . The term in the parentheses represents the coadded E -mode map, with observation weights w_i . Observation j is removed from the coadd to eliminate noise bias, since noise is assumed to be uncorrelated between observations.

The estimated true spectra (\hat{C}_b) are calculated from the pseudospectra using

$$\hat{C}_{b,j}^{EB} \equiv \sum_{b'} K_{bb',j}^{-1} P_{b',j} \tilde{C}_{\ell,j}^{EB} \quad (10)$$

and

$$K_{bb',j} = \sum_{\ell\ell'} P_{b\ell} M_{\ell\ell',\text{patch}} F_{\ell',j} B_{\ell'}^2 Q_{\ell'b'}. \quad (11)$$

All variables in this calculation have the same meaning as in PB17 but some are re-calculated for this analysis. The mode-mixing matrices $M_{\ell\ell',\text{patch}}$ correct for the effects of apodization and are calculated analytically for each patch using the patch apodization mask, discussed in Sec. III C. The filter transfer functions $F_{\ell',j}$ correct for the effects of time-domain filters for observation j , B_ℓ corrects for the beam, and $P_{b\ell}$ and $Q_{\ell'b'}$ are binning matrices. There are

four bins b centered at $\ell = [700, 1100, 1500, 1900]$, each of width $\Delta\ell = 400$.

E. Spectra to angles

Since the noise between different observations is uncorrelated, the noise between different C_ℓ^{EB} spectra is also uncorrelated. We will neglect the impact of foregrounds in the angle estimation: all three patches were chosen to have low foregrounds, and this systematic effect is discussed in Sec. V C. The estimated rotation angle $\hat{\alpha}_j$ for observation j is then obtained by minimizing

$$\chi^2(\hat{\alpha}_j) = \sum_{bb'} (\hat{C}_{b,j}^{EB} - 2\hat{\alpha}_j C_b^{EE,\text{th}}) (V_{bb'}^{EB})^{-1} (\hat{C}_{b',j}^{EB} - 2\hat{\alpha}_j C_{b'}^{EE,\text{th}}). \quad (12)$$

The theoretical power spectrum $C_{b'}^{EE,\text{th}}$ is calculated from the WMAP 9-year data. We use WMAP rather than Planck data because there are negligible differences for the purpose of our analysis and the WMAP results are already integrated in our pipeline. The covariance matrix $V_{bb'}^{EB}$ is calculated from 500 simulated maps that contain Λ CDM signal and instrumental noise. These simulated maps are generated by constructing maps containing only Λ CDM signal and scanning them into TOD. White noise from the PB17 noise model is added before running the maps through the full analysis pipeline. These simulations are also used to generate an error term σ_j for each observation, which is the standard deviation of all simulated angles for that observation. This construction includes the effect of sample variance, which is not relevant for our analysis because we make many observations of the same fixed CMB realization. The effect is negligible, however, because the angle error from sample variance alone is at least $16\times$ smaller than the total statistical error for each observation. The set of angles used in this analysis is shown in Fig. 1.

As in PB17, the absolute polarization angle of the instrument is calibrated using the EB -derived angle

coadded across both seasons. This means that the weighted mean angle in this analysis is fixed to zero. While other non-*EB* calibration sources exist, this procedure, described in [44], is the most accurate way to calibrate the instrument. We show in Sec. III F 2 that this absolute polarization angle calibration has a negligible effect on our analysis.

The TOD filtering causes *E*-mode power to leak into *B*-mode maps, which can be estimated via simulation. 100 simulations of maps containing only CMB C_ℓ^{TT} and C_ℓ^{EE} power are run through the pipeline and used to construct $\hat{C}_{b,j}^{EB}$. A small amount of leakage is detected. It is less than 6% of the statistical error for any angle, resulting in a negligible $< 0.01^\circ$ bias at any axion frequency probed in this analysis, and so we ignore its impact.

F. Likelihood

The noise between observations is assumed to be uncorrelated, and the errors on each angle Gaussian. The likelihood for the sinusoid parameters is then

$$\mathcal{L}(A, f, \theta) \propto e^{-\frac{1}{2}\chi^2(A, f, \theta)} \quad (13)$$

where

$$\chi^2(A, f, \theta) = \sum_j^N \frac{(\langle s_j \rangle - \hat{\alpha}_j)^2}{\sigma_j^2}. \quad (14)$$

The index j represents a single observation, and the signal $\langle s_j \rangle$ is defined below. The likelihood normalization factor $\prod_j^N \sigma_j \sqrt{2\pi}$ is dropped because we are only interested in relative likelihoods. Observations which have σ_j at least two times larger than the mean value are discarded, since they contribute negligibly to the likelihood and are more likely to have systematic errors. This removed 23 observations, which constituted $< 0.5\%$ of the statistical weight of the dataset. This leaves $N = 515$ angles available for the analysis.

The signal $\langle s_j \rangle$ must be averaged over the duration of each observation:

$$\begin{aligned} \langle s_j \rangle &= \frac{1}{t_j^{\text{end}} - t_j^{\text{start}}} \int_{t_j^{\text{start}}}^{t_j^{\text{end}}} dt A \sin(2\pi f t + \theta) \\ &= \text{Asinc}(\pi f \Delta t_j) \sin(2\pi f \bar{t}_j + \theta). \end{aligned} \quad (15)$$

$\Delta t_j \equiv t_j^{\text{end}} - t_j^{\text{start}}$ is the observation duration and $\bar{t}_j \equiv (t_j^{\text{end}} + t_j^{\text{start}})/2$ is the middle of the observation time. The effect of averaging is to wash out signals with oscillation periods that are shorter than the 4–8 hour observation periods.

1. Frequency domain

The frequency domain for our search is

$$0.02 \text{ days}^{-1} \leq f \leq 0.45 \text{ days}^{-1}. \quad (16)$$

The dataset spans 694 days, so we are able to probe sinusoidal periods out to approximately that length. However, as explained in Sec. V, uncertainty in the relative HWP angle between the two seasons may show up as a low-frequency signal, leading us to set a minimum frequency of 0.02 days^{-1} . The maximum frequency is set to 0.45 days^{-1} because the typical time-spacing between observations of a single patch is one day. This causes a significant loss in sensitivity around a frequency of 0.5 days^{-1} , as well as integer multiples of this frequency. Furthermore, at frequencies higher than 0.5 days^{-1} the averaging effect in Eq. (15) reduces the amplitude of the signal beyond the level of a few percent. For these reasons, we choose to only analyze frequencies up to 0.45 days^{-1} .

We assess the likelihood at a set of 1,492 evenly-spaced frequencies. In the absence of noise, this allows the likelihood estimator to recover an injected signal at any frequency to within 5% of the correct amplitude. This results in a frequency resolution about 5 times smaller than $1/694 \text{ days}^{-1}$, which is the spacing that a discrete Fourier transform of the same length of data would have. As a result, likelihood results from neighboring frequencies are highly correlated. This does not pose an issue for our analysis because we compare the results to simulated distributions which do not assume independence between different frequencies.

2. Effect of absolute polarization angle calibration

As described in Sec. III E, the absolute polarization angle of the instrument is calibrated by setting the weighted mean angle to zero. This means that our analysis is not sensitive to a static birefringence angle. To first order this is not an issue because we are searching for an oscillatory signal. However in the presence of such a signal, there will generally be some small shift in the mean angle caused by a partial oscillation. This shift is removed by the mean angle subtraction. This effect only becomes significant when there is $\lesssim 1$ oscillation period in the entire dataset. Our maximum period is 50 days, which corresponds to about 14 oscillation periods over the 694-day duration of our dataset. Therefore the subtraction in mean angle has negligible impact on the recovery of the signal amplitude.

IV. NULL TESTS

We run a series of null tests to check for systematic errors in our analysis before unblinding the data. The data are split along 15 possible sources of potential contamination, differenced to cancel any true signal, and tested for the presence of systematics. Twelve of the splits are the same as

in PB17: “moon distance,” “sun distance,” “rising vs. setting,” “high vs. low precipitable water vapor (PWV),” “high vs. low elevation,” “high vs. low gain,” “sun above or below horizon,” “1st vs. 2nd season,” “1st vs. 2nd half of dataset,” “left- vs. right-going scans,” “left vs. right side of the focal plane,” and “pixel polarization type.” We also add three new patch null tests, which pair two patches together in all three possible ways. Each split in these tests contains only the angles that were made from the B -mode maps from that patch, allowing us to probe for differences between patches.

A. Methodology

Two sets of angles are constructed for each test, one per split. The angles are constructed following the procedure in Sec. III, with the following modifications:

- (i) The B -mode maps contain only data from the split. This means that each map may contain less data or be eliminated altogether.
- (ii) For each split, a new map apodization mask is constructed per patch. These masks reflect differences in map coverage for each split. There are a few observations that have excessive noise in the unmasked region of some splits, and therefore cannot be null-validated. These observations’ B -mode and E -mode maps are excluded from the analysis in this paper, as described in Sec. III C.

For each split, *all* of the E -mode map data is used to construct C_e^{EB} , not just data from the split. This is because we wish to test for time-dependent systematic issues with the B -mode maps, and consider the coadded E -mode map validated in previous analyses (PB17).

A set of null statistics are computed for each test. For each of the two sets of angles “1” and “2,” at each frequency f , we find the maximum likelihood estimation (mle) of the amplitude A^{mle} and phase θ^{mle} that maximizes $\mathcal{L}(A, f, \theta)$. The null statistic at a given frequency is

$$T_{\text{null}}(f) \equiv \frac{|A_{\text{null}}(f)|^2}{\sigma(\Re(A_{\text{null}}(f)))^2} \quad (17)$$

where

$$A_{\text{null}}(f) \equiv (A_1^{\text{mle}} e^{i\theta_1^{\text{mle}}} - A_2^{\text{mle}} e^{i\theta_2^{\text{mle}}})(f). \quad (18)$$

The T_{null} can be thought of as a generalization of the difference of two discrete Fourier transforms: the correspondence would be exact if we had instantaneous observations evenly spaced in time and equally weighted. The normalization factor $\sigma(\Re(A_{\text{null}}))$ is calculated using 500 simulations and ensures that $T_{\text{null}}(f)$ is distributed approximately as chi-squared with 2 degrees of freedom. This is because $\Re(A_{\text{null}})$ and $\Im(A_{\text{null}})$ are both distributed approximately as Gaussians with the same variance and zero mean. We do not rely on these analytic distributions since

TABLE I. The five null test PTE values used in the pass criteria #1.

PTE statistic	Description	PTE
$\max_{t,f} T_{\text{null}}$	Spurious axion signal	0.032
$\sum_{t,f} T_{\text{null}}$	Total chi-square	0.062
$\max_t \sum_f T_{\text{null}}$	Bad test	0.060
$\max_f \sum_t T_{\text{null}}$	Bad frequency	0.246
$\max_t T_{\text{null}}(f=0)$	Mean angle offset	0.192

we compare the results to simulations, but we do rely on the normalization to allow for fair comparison across different tests and frequencies.

We calculate T_{null} at $f=0$ for the 10 tests out of 15 that have associated periods within the frequency range of this search. This statistic is simply the difference in the weighted mean angles between the two splits. We include it because differences in mean angle between two splits could cause a spurious signal with a period associated with the split: for example, a difference in angle caused by the moon distance could generate a signal aligned with the ~ 28 days synodic cycle. Five tests are omitted from the $f=0$ analysis. “1st vs. 2nd season” and “1st vs. 2nd half of dataset” are excluded because the associated period is much longer than our maximum period of 50 days. The “left- vs. right-going scans,” “left vs. right side of the focal plane,” and “pixel polarization type” are excluded because both splits contain all observations, so a difference in mean angle would not generate a time-dependent signal. In total, there are 22390 T_{null} values: 1492 frequencies from Sec. III F 1 for each of the 15 tests, and $f=0$ for 10 tests.

B. Pass criteria

There are two pass criteria that need to be satisfied in order to pass the null tests. The first, “pass criteria #1,” assesses failure in individual parts of the null test suite. The second, “pass criteria #2,” uses Kolmogorov-Smirnov (KS) tests to assess uniformity of all null suite probabilities.

Using T_{null} , five probability-to-exceed (PTE) values are computed. The lowest PTE among the five, $P_{\text{low}}^{\#1}$, is compared to a distribution generated from 500 simulations to calculate a global significance value, $\text{PTE}(P_{\text{low}}^{\#1})$. Pass criteria #1 requires more than 5% of simulations to have a lower “lowest of five PTEs” value than $P_{\text{low}}^{\#1}$: i.e. $1 - \text{PTE}(P_{\text{low}}^{\#1}) > 0.05$.

The five PTE values are shown in Table I, and are

- (1) $\max_{t,f} T_{\text{null}}$: The maximum T_{null} value across all tests and nonzero frequencies. This tests for the presence of a systematic sinusoidal signal.
- (2) $\sum_{t,f} T_{\text{null}}$: The total chi-square value over all tests and nonzero frequencies, which assesses noise mis-estimation.

TABLE II. The three null test PTE values used in the pass criteria #2.

Axion KS Test inputs	Description	Number of inputs	PTE
$\text{PTE}_{f,t}(T_{\text{null}})$	Overall	22380	0.128
$\text{PTE}_f(\sum_t T_{\text{null}})$	Per frequency	1492	0.122
$\text{PTE}_t(\sum_f T_{\text{null}})$	Per test	15	0.190

- (3) $\max_t \sum_f T_{\text{null}}$: The maximum per-test total chi-square value, which assesses issues with an individual test.
- (4) $\max_f \sum_t T_{\text{null}}$: The maximum per-frequency total chi-square value, which assesses issues with an individual frequency.
- (5) $\max_t T_{\text{null}}(f=0)$. The maximum difference in mean angle offset over the ten $f=0$ null tests.

The first four PTEs are calculated using only nonzero frequencies.

Three PTE values are computed using KS tests, and pass criteria #2 is satisfied if more than 5% of simulations have a lower value than $P_{\text{low}}^{\#2}$: i.e. $1 - \text{PTE}(P_{\text{low}}^{\#2}) > 0.05$. Here $P_{\text{low}}^{\#2}$ is the lowest of the three PTEs. These tests assess whether the probabilities of getting various T_{null} values follow a uniform distribution. The KS test inputs and PTE results are shown in Table II. The probability value returned by the KS test cannot be used to assess significance, since correlations between null tests result in a slightly non-uniform distribution of individual probabilities. Instead, the value of the KS statistic itself is compared to simulations.

C. Null test results

For pass criteria #1, $1 - \text{PTE}(P_{\text{low}}^{\#1}) = 0.124$. For pass criteria #2, $1 - \text{PTE}(P_{\text{low}}^{\#2}) = 0.148$. Both are greater than 0.05, and so the null tests pass.

While the null tests pass our stated criteria, all of the PTE values are relatively low. This is potentially indicative of a systematic issue, which we investigated. The low PTEs do not appear to be caused by particular null splits or frequency ranges. Two of the PTEs in Table I (“bad test” and “bad frequency”) explicitly look for the worst test and frequency, and were not discrepant enough to fail pass criteria #1. Pass criteria #2 is sensitive to non-uniformities in the T_{null} distribution and also passed. In addition to these summary statistics, visual inspection of the full set of T_{null} values plotted for each test and frequency do not reveal large asymmetries between different null splits and frequency ranges. A set of correlated, low PTEs is not surprising given a high total chi-square value. The real data has a total chi-square value 20% larger than the simulation average. This is larger than all but 6% of simulations as shown in Table I. Simulations with such

a large total chi-square value have the overwhelming majority of their PTE values less than 0.2.

A high total chi-square value could indicate a problem with the noise model. The noise model is based on Monte-Carlo simulations of white noise timestreams, as described in PB17. It was validated on the full coadded spectra, so we do not expect it to be a source of error. However, to check for noise misestimation, we computed a set of angles using a different 100 realizations of sign flip noise B -mode maps, which are created by randomly reversing the sign of the maps made for each 15 minute CES. This cancels out the true signal but maintains the noise properties of the real data. We compared $\sigma(\Re(A_{\text{null}}))$ from sign flip simulations to those computed using Monte-Carlo noise simulations, and found they agreed to within 5%, and in the wrong direction to alleviate the total chi-square tension. Having found no obvious discrepancy under a different noise model, and with the null tests having passed our pre-determined criteria, we decided that a somewhat elevated total chi-square value would not prevent us from unblinding the data.

V. SYSTEMATIC ERRORS

In this section we examine the impact of several sources of systematic error, including uncertainty in the position of the HWP, any differences in average angle between the three observing patches, and foregrounds. Although we rely on some of the systematic pipeline used in PB17, our analysis is fundamentally different because we are looking for *time-dependent* sources of error. Systematics that provide a constant C_{ℓ}^{EB} offset are irrelevant. Therefore many of the PB17 systematics are of no concern, while two others, the HWP and differences between patches, require careful consideration. The typical maximum likelihood sinusoid amplitude generated by statistical noise alone is 0.26° . Therefore any oscillatory systematic that generates an amplitude much less than 0.26° should have negligible impact on the limits in this work.

A. HWP position uncertainty

As discussed in Sec. II, angle errors that are introduced by the stepped rotation of the HWP during the first observing season are a significant source of time-dependent uncertainty. The HWP angle was changed about 60 times during the first half of the first season, 4 times during the second half of the first season, and was then fixed for the second season. During each rotation the HWP was commanded to rotate in increments of 11.25° . However, when examining the polarization angles derived from an alternative calibration source, the Crab Nebula (Tau A), it became clear that the angles during the first season exhibited larger variance than during the second season. Furthermore, there was an offset in angle between the two seasons. Both phenomena could be explained if the HWP




FIG. 2. The estimated typical maximum likelihood amplitude signal of the HWP angle offsets. 500 simulations containing only random offset angles were generated and run through the likelihood. The median recovered amplitude is shown here. The HWP step cadence begins to generate larger systematic issues at periods > 50 days. Based on these results, the maximum period used in this analysis is 50 days.

did not step exactly to the commanded position each time. While the exact HWP-induced offset at each step is unknown due to statistical error on the Tau A measurements, PB17 found that the discrepancy could be explained by adding a systematic error of 0.56° for each step, corresponding to a typical HWP offset of 0.28° , in quadrature with the statistical error.

This systematic is highly important for this analysis because it corresponds to a unknown time-dependent shift in the instrument polarization angle, which could mimic an axion-generated birefringence signal. The HWP-induced offset between the first and second seasons, for example, is degenerate with a signal with a period of one year. To assess the impact at higher frequencies, we ran 500 simulations where a random angle offset, drawn from a Gaussian distribution with $\sigma = 0.56^\circ$, was assigned to each HWP angle. Then we generated timestreams where each observation’s CMB angle corresponded to the offset given by the HWP angle. The maximum likelihood sinusoid amplitude was then estimated at each axion frequency.

The result is shown in Fig. 2. At high frequencies the HWP rotation cadence does not correspond to any specific frequency. Since observations have a statistical error of at least 2° and the typical HWP offset is 0.56° , the impact at high frequencies is minimal. On average these offsets cause only a 1%–2% increase in the maximum likelihood amplitude estimate at a given frequency relative to statistical noise alone, and so we ignore the impact of this HWP offset noise in our likelihood. At low frequencies, however, the last few HWP steps in the first season and the offset between the first and second seasons begin to translate into larger sinusoidal signals. We emphasize that because the

exact HWP-induced offset at each step is unknown, the numbers shown in Fig. 2 are only indicative of the average effect we expect, not the true effect present in our data. Because the HWP step cadence begins to generate larger systematic issues at periods > 50 days, there is a higher potential to generate a signal that would cause false detection. Based on these results, we chose to set 50 days as the longest period we analyze.

It is possible that the HWP offsets follow a more pathological model than what we assume and mimic a sinusoidal signal at some frequency in our domain. The “1st half vs 2nd half” null test described in Sec. IV tests for this, because the time periods align with when the HWP was stepped vs. fixed. As stated in that section, there is no signal observed in this null test.

B. Differences in patch mean angle

Since a typical day involves observing all three patches sequentially, any differences in the mean angle of each of the three patches could generate a oscillatory signal with a period of one day. However, as explained in Sec. III F 1, our upper frequency bound is 0.45 days^{-1} and so we are not sensitive to this signal. However, variations in the observing sequence can generate signals at other frequencies, and simulations of patch mean angle offsets reveal a transfer of $\leq 12\%$ of any offset into any frequency in domain we do observe. Fortunately, our null tests directly calculate the difference in mean angle between patches. Since they are consistent with simulations, which do not contain foreground power, we do not see evidence for any foreground or systematic patch mean angle differences.

One known effect that will cause offsets between the patch mean angles is sample variance, because each patch has an independent sample-variance-induced C_ℓ^{EB} signal. This effect is included in the Λ CDM simulations, but if it were large we would need to correct for its effects. It can be calculated both analytically and through simulations, and in each patch, the error from sample variance is $\sigma \approx 0.12^\circ$. This is $3 \times$ smaller than the statistical error alone on each patch mean angle. Furthermore, since the patch mean angle offsets transfer $\leq 12\%$ into any given frequency, the typical A^{mle} generated by cosmic variance is a negligible $\lesssim 0.01^\circ$.

C. Foregrounds and other instrumental systematics

Any C_ℓ^{EB} contribution from dust and synchrotron radiation should be time-independent and simply contribute to the patch mean angle differences. These differences were addressed in Sec. V B. Point sources may be variable in time, but the strongest are masked as discussed in PB17.

PB17 also makes many other estimates related to the calibration, analysis effects, and known instrumental systematics. We expect the multiplicative effects to have a similar impact on the C_ℓ^{EB} spectrum as they do on the C_ℓ^{BB} spectrum, which PB17 determined was 6%. Since the angle

estimates are linear in C_ℓ^{EB} , we would then roughly expect a overall 6% uncertainty on amplitude estimation. Since this is relatively small and does not bias our results, this is negligible uncertainty in our reported results.

The other instrumental systematics deal with effects that could potentially cause spurious additive C_ℓ^{EB} , including differential gain, gain drift, differential beam size, differential ellipticity, differential pointing, boresight pointing, and electrical crosstalk. These can impact our analysis if they generate a time-dependent C_ℓ^{EB} . An analysis of C_ℓ^{EB} bias was done in PB14, which noted that all sources provided $\lesssim 0.02^\circ$ bias in the coadded spectra. Using a combination of analytic estimates [46] and simulations, we modified the analysis to account for per-observation C_ℓ^{EB} and determined that all such sources were either time-independent or provided a negligibly small time-dependent contribution as compared to the statistical noise.

VI. RESULTS

We search for the presence of a sinusoidal signal, and finding no statistically significant indication, place appropriate bounds over the frequency range of the search. By accounting for the axion field amplitude at the telescope, these translate to bounds on the axion-photon coupling over a range of axion masses.

A. Search for a signal

In order to detect a signal, we form a test statistic $\Delta\chi^2$ which is large in the presence of a sinusoidal signal above the background. For each frequency in the discrete frequency domain, the MLE phase and amplitude are found, and the difference relative to the χ^2 with no signal is calculated:

$$\Delta\chi^2(f) \equiv \chi^2(A=0) - \chi^2(A^{\text{mle}}(f), f, \theta^{\text{mle}}(f)). \quad (19)$$

The test statistic is the largest $\Delta\chi^2(f)$ among all frequencies:

$$\Delta\chi^2 \equiv \max_f(\Delta\chi^2(f)). \quad (20)$$

This statistic is optimal in the Neyman-Pearson sense in that it maximizes the probability of rejecting the no-signal hypothesis if the alternative hypothesis of a sinusoidal signal is true [47]. The distribution of the statistic is approximately chi-squared, but we do not rely on this, instead comparing to a distribution of test statistics computed from 500 Λ CDM + noise simulations. The p-value is the fraction of background simulation test statistic values that exceed that of the real data.

The selection of $\Delta\chi^2$ and its location relative to the distribution is shown in Fig. 3. The real data has a p-value of 0.048, corresponding to 1.7σ significance. While this p-value is relatively small, it is not significant enough to claim detection of a signal above the noise background. Future analysis of the remaining three seasons of POLARBEAR




FIG. 3. The significance test for the presence of an axion signal. Top: the calculation of $\Delta\chi^2$. For each frequency in the discrete frequency domain, the MLE phase and amplitude are found, and $\Delta\chi^2(f)$ is calculated. $\Delta\chi^2$ is the largest among these. Bottom: the significance of $\Delta\chi^2$ as compared to a set of 500 simulations.

will aid in determining if this result is simply a statistical fluctuation.

When assessing the significance of a result with many trials, in this case discrete frequencies, it is often useful to report the experimental sensitivity. This is defined as the signal amplitude which would produce a detection at some significance level. We can approximate the $\Delta\chi^2$ required for any significance level by applying a fit to the distribution shown in the bottom of Fig. 3. A 3σ detection (PTE = 0.0013) would require $\Delta\chi^2 = 28.5$, yielding an experimental sensitivity of $A(3\sigma) = 1.15^\circ$. The 5σ (PTE = 2.8×10^{-7}) experimental sensitivity is $A(5\sigma) = 1.43^\circ$.

B. Upper limit

As we do not detect a globally significant axion-like oscillation, at each frequency we compute a frequentist

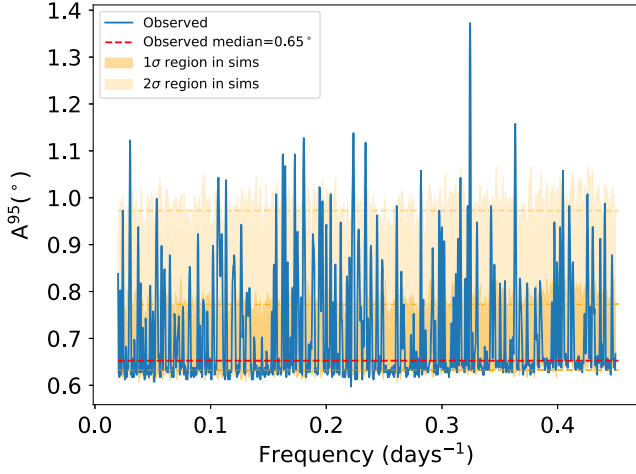


FIG. 4. A 95% upper confidence limit on the presence of a sinusoidal signal at each frequency. The 1σ and 2σ regions contain approximately 68% and 95% of the upper bounds calculated from 500 simulations at each frequency. The median values for these regions are indicated with dashed lines. Many individual frequencies exceed the 2σ threshold, which is to be expected based on the large number of test frequencies, and does not necessarily indicate a detection.

upper limit on the signal amplitude at 95% confidence level. The likelihood for the amplitude as a function of frequency is

$$\mathcal{L}_A(A, f) = \int_0^{2\pi} d\theta \mathcal{L}(A, f, \theta) P(\theta), \quad (21)$$

with a uniform probability distribution for the phase

$$P(\theta) = \frac{1}{2\pi}, \quad 0 \leq \theta < 2\pi. \quad (22)$$

We construct confidence intervals following the Neyman procedure [47] with the following ordering: intervals $[0, A^{95}(f)]$ are defined with an upper limit $A^{95}(f)$ at each frequency f such that

$$P(A^{\text{mle}} \leq A_{\text{obs}}^{\text{mle}} | A^{95})(f) = 0.05. \quad (23)$$

Here $A_{\text{obs}}^{\text{mle}}$ is the observed maximum likelihood amplitude calculated from the data using \mathcal{L}_A . $P(A^{\text{mle}} | A)(f)$ is the probability of calculating A^{mle} given a true signal amplitude A plus background noise. Therefore $P(A_{\text{mle}} \leq A_{\text{obs}}^{\text{mle}} | A)(f)$ represents the probability that, given a true signal of amplitude A , A^{mle} would be less than or equal to the value observed. This probability is a monotonically decreasing function of A , and so it is less than 0.05 for all A -values above the upper limit. In other words, signal amplitudes that would generate an A^{mle} as low or lower less than 5% of the time are excluded from the confidence interval. The

limits calculated by this procedure are shown in Fig. 4. The median upper bound is 0.65° . The bound varies over frequencies as expected, in analogy to the typical behavior of a Fourier transform.

The probability distribution $P(A^{\text{mle}} | A)(f)$ is generated by calculating A^{mle} from 500 simulated angle timestreams of background noise with an injected signal of amplitude A , frequency f , and random phase between 0 and 2π . This is done for all frequencies and discretized array of A -values, which are smoothed to create a continuous probability distribution. The results are checked for convergence of the median upper limit to the $< 1\%$ level. $P(A^{\text{mle}} | A)(f)$ is almost the same for each frequency. The primary difference comes from the finite duration of each observation, which reduces the strength of the bounds at higher frequencies according to the sinc function in Eq. (15). We can approximate the impact of this by calculating $\text{sinc}(\pi f \overline{\Delta t})$, where $\overline{\Delta t} = 6.4$ hours is the weighted mean observation duration. The maximum amplitude reduction occurs at the largest frequency and is $< 2.5\%$. While this effect is included in the constraint we place, it is small enough that we still report the median bound over the full frequency range.

C. Constraints

The median angle bound A^{95} represents a search for a sinusoidal signal in the CMB data independent of any axion model. Translating this into a constraint on the axion-photon coupling constant $g_{\phi\gamma}$ requires specifying model parameters. In F19, BK22, and SPT22, this was done by simply requiring that the local axion energy density equal the average Milky Way dark matter energy density, which we will call the “deterministic” case, following Ref. [22]. This means the axion field has amplitude ϕ_{DM} , where $\phi_{\text{DM}}^2 m_\phi^2 / 2 = \rho_0$ and ρ_0 is the local density of dark matter. This gives the relation

$$g_{\phi\gamma} = (1.6 \times 10^{-11} \text{ GeV}^{-1}) \times \left(\frac{A}{1^\circ} \right) \times \left(\frac{m_\phi}{10^{-21} \text{ eV}} \right) \times \left(\frac{\kappa \times \rho_0}{0.3 \text{ GeV/cm}^3} \right)^{-1/2}. \quad (24)$$

Here A is the rotation amplitude, m_ϕ is the axion mass, and κ is the fraction of the dark matter that the axion constitutes. Using $\kappa = 1$, $\rho_0 = 0.3 \text{ GeV/cm}^3$, and $A^{95} = 0.65^\circ$, this constraint is

$$g_{\phi\gamma} < (1.1 \times 10^{-11} \text{ GeV}^{-1}) \times \left(\frac{m_\phi}{10^{-21} \text{ eV}} \right). \quad (25)$$

Reference [22], however, has recently pointed out that this constraint is inappropriate when the timescale of the experiment is much less than the axion field coherence




FIG. 5. Bounds on the axion-photon coupling as a function of axion mass. Shown in solid red is the 95% upper limit obtained in this work using a stochastic local axion field amplitude and assuming the axion constitutes all the dark matter (see Sec. VI). The lighter oscillating result is the exact bound, and the median result, $g_{\phi\gamma} < (2.4 \times 10^{-11} \text{ GeV}^{-1}) \times (m_{\phi}/10^{-21} \text{ eV})$, is shown in darker red. The median deterministic bounds for POLARBEAR, BICEP, and SPT are the dashed lines: these are the published results in [19] and [20]. The green “washout” bound was calculated in F19 from the lack of CMB polarization suppression. Lower bounds on the axion mass come from Milky Way satellites [33] and the Lyman- α forest [31]. Upper bounds on the axion-photon coupling from CAST [35] are also shown.

time, as it is in our case. When considering the local axion field as the sum of many individual wave modes, each with random phase, the amplitude ϕ_0 at any given time is a Rayleigh-distributed stochastic variable centered on ϕ_{DM} that varies with time on the coherence timescale $\tau_{\text{coherence}} \sim (fv_{\text{MW}}^2/c^2)^{-1}$, where the virial velocity is $v_{\text{MW}} \approx 10^{-3}c$. We will call this the “stochastic” case. This random phase model has been shown to roughly agree with simulations of fuzzy dark matter [10]. The largest frequency we consider in our analysis is $f_{\text{max}} = 0.45 \text{ days}^{-1}$, yielding a minimum coherence time of about 6000 years: therefore, for the purposes of our experiment we can approximate ϕ_0 as a fixed random variable. The phase is treated in the same manner with a uniform probability distribution. To form the likelihood for $g_{\phi\gamma}$ we integrate over these two parameters:

$$\mathcal{L}_{g_{\phi\gamma}}(g_{\phi\gamma}, f) = \int_0^{\infty} d\phi_0 \int_0^{2\pi} d\theta \mathcal{L}(g_{\phi\gamma}\phi_0/2, f, \theta) P(\phi_0) P(\theta). \quad (26)$$

Here $P(\phi_0)$ is a Rayleigh distribution centered on ϕ_{DM} :

$$P(\phi_0) = \frac{2\phi_0}{\phi_{\text{DM}}^2} e^{-\frac{\phi_0^2}{\phi_{\text{DM}}^2}}. \quad (27)$$

The resulting median 95% upper limit in this stochastic case is calculated in the same manner as Sec. VIB, and is

$$g_{\phi\gamma} < (2.4 \times 10^{-11} \text{ GeV}^{-1}) \times \left(\frac{m_{\phi}}{10^{-21} \text{ eV}} \right). \quad (28)$$

This is $2.2\times$ larger than the deterministic case.¹ The increase is due to the possibility that we happen to be observing at an unlucky time when $\phi_0 < \phi_{\text{DM}}$, which generates a smaller signal amplitude at a given $g_{\phi\gamma}$.

We briefly comment on the choice between frequentist and Bayesian statistics in this analysis. In the frequentist approach adopted here, we generate confidence intervals according to Eq. (23). In a Bayesian approach, we would instead generate a posterior probability distribution for $g_{\phi\gamma}$, which requires choosing a prior $P(g_{\phi\gamma})$.

In the deterministic case, or equivalently when placing limits on the sinusoid amplitude A , the results roughly agree for several choices of prior. With a uniform prior on A , the median Bayesian upper limit is 0.60° . If we had instead chosen to parametrize the signal [Eq. (3)] as

$$\beta_{\text{CMB}}(t) = B \sin(2\pi ft) + C \cos(2\pi ft), \quad (29)$$

applying uniform priors on B and C give a median upper limit of 0.71° . Both of these results are close to the median frequentist 0.65° upper limit we report in Sec. VI B.

In the stochastic case there is a large dependence on the choice of prior and parameterization. We must integrate over $P(\phi_0)$, and the nonzero probability of small $\phi_0 < \phi_{\text{DM}}$ values means that the posterior for $g_{\phi\gamma}$ has a long tail at large $g_{\phi\gamma}$. A uniform prior on $g_{\phi\gamma}$ results in a Bayesian upper limit $10\times$ greater than the deterministic case, much larger than the frequentist result. Furthermore, applying uniform priors in the sine-cosine parametrization [Eq. (29)] causes the upper bound to diverge. To resolve these issues, Ref. [22] advocates choosing a Berger-Bernardo prior for $g_{\phi\gamma}$, which like the Jeffreys's prior is parameterization-invariant, and in their analysis agrees with the frequentist result to within 3%. Unfortunately, for our likelihood we found no simple analytic form for this prior and it is difficult to accurately compute numerically. Due to issues with defining a Bayesian prior in the stochastic case, we choose to report a frequentist limit.

This limit is shown in Fig. 5, along with a selection of other constraints. Our primary result is the constraint on $g_{\phi\gamma}$ assuming that the axion field amplitude is a Rayleigh-distributed stochastic variable. It is shown in full detail along with the median limit from Eq. (28). The median deterministic constraint, Eq. (25), is also shown. The published results from BK22 and SPT22 are deterministic Bayesian upper bounds with a uniform prior on $P(g_{\phi\gamma})$. None of the other bounds shown in Fig. 5 come from

¹Reference [22] reports an increase of $2.7\times$ when using the frequentist approach. The difference possibly stems from the use of a different likelihood (they assume that the data is uniformly spaced in time) and/or differences between individual realizations of the noise. When comparing the Bayesian deterministic approach and the Bayesian stochastic approach with a uniform amplitude prior we see a $10\times$ increase, in agreement with [22].

assuming a value for the axion field in a dark matter halo. We emphasize that this is the first CMB analysis of this effect that we are aware of to include the local stochastic nature of ϕ_0 .

In both the stochastic and deterministic cases, the bounds apply over the frequency range presented in Eq. (16), which corresponds to the axion mass range

$$9.6 \times 10^{-22} \text{ eV} \leq m_\phi \leq 2.2 \times 10^{-20} \text{ eV}. \quad (30)$$

VII. CONCLUSION

We have used POLARBEAR data to search for a coherent, all-sky, sinusoidal oscillation of the CMB polarization angle in time. We do not detect such a signal, and place a median 95% upper limit of 0.65° on the sinusoid amplitude over oscillation frequencies between 0.02 days^{-1} and 0.45 days^{-1} . We use these results to constrain the coupling between electromagnetism and an axion, here defined as an ultralight pseudoscalar field, under the assumption that the axion constitutes all of the dark matter. The signal depends on the value of the axion field at the telescope, and under the assumption that the field amplitude is a Rayleigh-distributed stochastic variable, we set the limit $g_{\phi\gamma} < (2.4 \times 10^{-11} \text{ GeV}^{-1}) \times (m_\phi/10^{-21} \text{ eV})$ over the mass range $9.6 \times 10^{-22} \text{ eV} \leq m_\phi \leq 2.2 \times 10^{-20} \text{ eV}$.

Three additional seasons of POLARBEAR data have been collected in addition to the two seasons analyzed here, and we anticipate that analyzing them will improve our data volume by a factor of $\approx 2-3$, with a corresponding $\approx 60\%$ improvement in the constraints. This data will also possibly allow us to probe lower frequencies (smaller axion masses), because these seasons are not affected by the HWP position uncertainty that restricted the frequency range in this analysis. A promising avenue for placing constraints several times better than this with POLARBEAR data lies with measurements of Tau A, which was used as a polarization calibration source, and was precisely measured during the five observing seasons between 2012 and 2016. These measurements present a different challenge than the CMB because the axion field at the source needs to be carefully considered, whereas in our analysis the $\mathcal{O}(10^5)$ year duration of recombination allowed for ignoring the source term. Nonetheless, this additional data provides another avenue to search for the presence of axions using POLARBEAR data.

Several future CMB experiments, including the Simons Array [48], Simons Observatory [49], and CMB-S4 [50], should be able to perform a similar analysis with improved constraints. This analysis imposes no additional requirements on the design or operation of these experiments, it simply requires making many measurements of the CMB over an extended period of time. The sensitivity is not fundamentally limited by anything other than the precision of the polarization measurements, unlike the CMB washout

effect, which is limited by cosmic variance [17]. The unknown amplitude of the axion field at the telescope is the chief source of model-dependence, but it is well described by a probability distribution that can be treated statistically when placing bounds. Unlike many other astrophysical measurements, this analysis does not suffer from significant modeling uncertainty at the polarization source due to the well-understood nature of the CMB. In the rapidly growing field of axion searches, this should allow future CMB experiments to provide increasingly competitive measurements of the axion-photon coupling constant.

ACKNOWLEDGMENTS

The POLARBEAR project is funded by the National Science Foundation under grants No. AST-0618398 and No. AST-1212230. The analysis presented here was also supported by Moore Foundation Grant No. 4633, the Simons Foundation Grant No. 034079, and the Templeton Foundation Grant No. 58724. C. B. acknowledges support from the COSMOS project of the Italian Space Agency, and the INDARK Initiative of the INFN (web.infn.it/CSN4/IS/Linea5/InDark). The James Ax Observatory operates in the Parque Astronómico Atacama in Northern Chile under the auspices of the Comisión Nacional de Investigación Científica y

Tecnológica de Chile (CONICYT). Y. C. acknowledges the support from JSPS KAKENHI Grants No. 18K13558 and No. 21K03585. J. E. acknowledges funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (Grant agreement No. 101044073). G. F. acknowledges the support of the European Research Council under the Marie Skłodowska Curie actions through the Individual Global Fellowship No. 892401 PiCOGAMBAS. This work was supported by World Premier International Research Center Initiative (WPI), MEXT, Japan. The work of M. H. was supported by JSPS KAKENHI Grant No. JP22H04945. In Japan, this work was supported by JSPS KAKENHI grant No. 18H05539 and No. 19H00674. This work was supported in part by JSPS core-to-core program number JPJSCCA20200003. Work at L. B. N. L. is supported in part by the U.S. Department of Energy, Office of Science, Office of High Energy Physics, under Contract No. DE-AC02-05CH11231. H. N. acknowledges the support from the JSPS KAKENHI Grant No. JP26800125. C. R. acknowledges support from the Australian Research Council’s Discovery Projects scheme (No. DP210102386). S. T. acknowledges JSPS Overseas Research Fellowships and JSPS KAKENHI Grant No. JP14J01662. Support from the James B. Ax Family Foundation is acknowledged.

-
- [1] L. D. Duffy and K. van Bibber, *New J. Phys.* **11**, 105008 (2009).
 - [2] D. J. Marsh, *Phys. Rep.* **643**, 1 (2016).
 - [3] L. Hui, *Annu. Rev. Astron. Astrophys.* **59**, 247 (2021).
 - [4] R. D. Peccei and H. R. Quinn, *Phys. Rev. Lett.* **38**, 1440 (1977).
 - [5] R. D. Peccei and H. R. Quinn, *Phys. Rev. D* **16**, 1791 (1977).
 - [6] S. Weinberg, *Phys. Rev. Lett.* **40**, 223 (1978).
 - [7] F. Wilczek, *Phys. Rev. Lett.* **40**, 279 (1978).
 - [8] A. Arvanitaki, S. Dimopoulos, S. Dubovsky, N. Kaloper, and J. March-Russell, *Phys. Rev. D* **81**, 123530 (2010).
 - [9] P. Svrcek and E. Witten, *J. High Energy Phys.* **06** (2006) 051.
 - [10] L. Hui, J. P. Ostriker, S. Tremaine, and E. Witten, *Phys. Rev. D* **95**, 043541 (2017).
 - [11] W. Hu, R. Barkana, and A. Gruzinov, *Phys. Rev. Lett.* **85**, 1158 (2000).
 - [12] A. Caldwell, G. Dvali, B. Majorovits, A. Millar, G. Raffelt, J. Redondo, O. Reimann, F. Simon, and F. S. and, *Phys. Rev. Lett.* **118**, 091801 (2017).
 - [13] Y. Kahn, B. R. Safdi, and J. Thaler, *Phys. Rev. Lett.* **117**, 141801 (2016).
 - [14] I. G. Irastorza and J. Redondo, *Prog. Part. Nucl. Phys.* **102**, 89 (2018).
 - [15] D. Harari and P. Sikivie, *Phys. Lett. B* **289**, 67 (1992).
 - [16] S. M. Carroll, *Phys. Rev. Lett.* **81**, 3067 (1998).
 - [17] M. A. Fedderke, P. W. Graham, and S. Rajendran, *Phys. Rev. D* **100**, 015040 (2019).
 - [18] P. A. R. Ade *et al.* (BICEP/Keck Collaborations), *Phys. Rev. D* **103**, 042002 (2021).
 - [19] P. A. R. Ade *et al.* (BICEP/Keck Collaborations), *Phys. Rev. D* **105**, 022006 (2022).
 - [20] K. R. Ferguson *et al.* (SPT-3G Collaboration), *Phys. Rev. D* **106**, 042011 (2022).
 - [21] J. W. Foster, N. L. Rodd, and B. R. Safdi, *Phys. Rev. D* **97**, 123006 (2018).
 - [22] G. P. Centers, J. W. Blanchard, J. Conrad, N. L. Figueroa, A. Garcon, A. V. Gramolin, D. F. J. Kimball, M. Lawson, B. Pelsers, J. A. Smiga, A. O. Sushkov, A. Wickenbrock, D. Budker, and A. Derevianko, *Nat. Commun.* **12**, 7321 (2021).
 - [23] T. Fujita, R. Tazaki, and K. Toma, *Phys. Rev. Lett.* **122**, 191101 (2019).
 - [24] A. Caputo, L. Sberna, M. Frías, D. Blas, P. Pani, L. Shao, and W. Yan, *Phys. Rev. D* **100**, 063515 (2019).
 - [25] T. Liu, G. Smoot, and Y. Zhao, *Phys. Rev. D* **101**, 063012 (2020).
 - [26] M. Ivanov, Y. Kovalev, M. Lister, A. Panin, A. Pushkarev, T. Savolainen, and S. Troitsky, *J. Cosmol. Astropart. Phys.* **02** (2019) 059.
 - [27] A. Basu, J. Goswami, D. J. Schwarz, and Y. Urakawa, *Phys. Rev. Lett.* **126**, 191102 (2021).

- [28] G.-W. Yuan, Z.-Q. Xia, C. Tang, Y. Zhao, Y.-F. Cai, Y. Chen, J. Shu, and Q. Yuan, *J. Cosmol. Astropart. Phys.* **03** (2021) 018.
- [29] Y. Chen, Y. Liu, R.-S. Lu, Y. Mizuno, J. Shu, X. Xue, Q. Yuan, and Y. Zhao, *Nat. Astron.* **6**, 592 (2022).
- [30] A. Castillo, J. Martin-Camalich, J. Terol-Calvo, D. Blas, A. Caputo, R. T. G. Santos, L. Sberna, M. Peel, and J. A. Rubiño Martín, *J. Cosmol. Astropart. Phys.* **06** (2022) 014.
- [31] V. Iršič, M. Viel, M. G. Haehnelt, J. S. Bolton, and G. D. Becker, *Phys. Rev. Lett.* **119**, 031302 (2017).
- [32] K. K. Rogers and H. V. Peiris, *Phys. Rev. Lett.* **126**, 071302 (2021).
- [33] E. O. Nadler *et al.* (DES Collaboration), *Phys. Rev. Lett.* **126**, 091101 (2021).
- [34] N. Dalal and A. Kravtsov, *Phys. Rev. D* **106**, 063517 (2022).
- [35] CAST Collaboration, *Nat. Phys.* **13**, 584 (2017).
- [36] A. Payez, C. Evoli, T. Fischer, M. Giannotti, A. Mirizzi, and A. Ringwald, *J. Cosmol. Astropart. Phys.* **02** (2015) 006.
- [37] J. Sisk-Reynés, J. H. Matthews, C. S. Reynolds, H. R. Russell, R. N. Smith, and M. C. D. Marsh, *Mon. Not. R. Astron. Soc.* **510**, 1264 (2021).
- [38] M. A. Buen-Abad, J. Fan, and C. Sun, *J. High Energy Phys.* **02** (2020) 103.
- [39] P. A. R. Ade *et al.* (POLARBEAR Collaboration), *Astrophys. J.* **794**, 171 (2014).
- [40] P. A. R. Ade *et al.* (POLARBEAR Collaboration), *Astrophys. J.* **848**, 121 (2017).
- [41] P. A. Ade *et al.* (POLARBEAR Collaboration), *Phys. Rev. D* **92**, 123509 (2015).
- [42] K. Arnold *et al.*, *Proc. SPIE Int. Soc. Opt. Eng.* **8452**, 84521D (2012).
- [43] Z. D. Kermish *et al.*, in *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VI*, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 8452, edited by W. S. Holland and J. Zmuidzinas (2012), p. 84521C, [arXiv:1210.7768](https://arxiv.org/abs/1210.7768).
- [44] B. G. Keating, M. Shimon, and A. P. S. Yadav, *Astrophys. J.* **762**, L23 (2012).
- [45] K. M. Smith, *Phys. Rev. D* **74**, 083002 (2006).
- [46] M. Shimon, B. Keating, N. Ponthieu, and E. Hivon, *Phys. Rev. D* **77**, 083003 (2008).
- [47] R. L. Workman *et al.* (Particle Data Group), *Prog. Theor. Exp. Phys.* **2022**, 083C01 (2022).
- [48] A. Suzuki *et al.*, *J. Low Temp. Phys.* **184**, 805 (2016).
- [49] P. Ade *et al.*, *J. Cosmol. Astropart. Phys.* **02** (2019) 056.
- [50] K. N. Abazajian *et al.*, *CMB-S4 Science Book*, first edition (2016).