

First indirect detection constraints on axions in the Solar basin

William DeRocco,^{1,*} Shalma Wegsman,^{2,†} Brian Grefenstette,^{3,‡} Junwu Huang,^{4,§} and Ken Van Tilburg^{2,5,¶}

¹*Santa Cruz Institute for Particle Physics, University of California, Santa Cruz, CA 95062, USA*

²*Center for Cosmology and Particle Physics, Department of Physics,
New York University, New York, NY 10003, USA*

³*Cahill Center for Astronomy and Astrophysics, California Institute of Technology, Pasadena, CA 91125, USA*

⁴*Perimeter Institute for Theoretical Physics, Waterloo, Ontario N2L 2Y5, Canada*

⁵*Center for Computational Astrophysics, Flatiron Institute, New York, NY 10010, USA*

(Dated: May 13, 2022)

Axions with masses of order keV can be produced in great abundance within the Solar core. The majority of Sun-produced axions escape to infinity, but a small fraction of the flux is produced with speeds below the escape velocity. Over time, this process populates a basin of slow-moving axions trapped on bound orbits. These axions can decay to two photons, yielding an observable signature. We place the first limits on this solar basin of axions using recent quiescent solar observations made by the NuSTAR X-ray telescope. We compare three different methodologies for setting constraints, and obtain world-leading limits for axions with masses between 5 and 30 keV, in some cases improving on stellar cooling bounds by more than an order of magnitude in coupling.

I. INTRODUCTION



The Sun is an exquisite laboratory for new physics. Studies of particles produced in the Sun have provided us with a wealth of information about light, weakly-interacting particles, including neutrinos [1–3], axions [4–6], and dark photons [7–9]. Most of these studies have focused solely on the flux of new particles that escape the gravitational potential of the Sun. However, a tiny fraction will be produced at sufficiently small velocities to become gravitationally trapped, forming a “Solar basin” of new particles. Ref. [10] showed that despite a small injection rate, the local density in this basin can grow to an appreciable fraction of the local dark matter density, thus providing a new target for direct detection experiments.

If these particles are unstable, however, their density in the Solar basin ceases to grow when decays balance injections, roughly when the Sun’s age is comparable to the particle’s lifetime. Once this steady-state is reached, the basin decays yield a constant flux of decay products. In the case of keV-mass axion-like particles [11–15] (hereafter “axions”) with a coupling to photons, these decays appear as a narrow line in the X-ray band with a characteristic angular distribution. (Certain aspects of these phenomena were foreseen in Ref. [16], which attempted to address the solar coronal heating problem with a Kaluza-Klein spectrum of gravitationally-trapped axions.)

This signal is therefore best probed by X-ray telescopes capable of direct solar observations, such as NuSTAR [17]. While not initially designed as a solar ob-

servatory, it has made observations of both the active Sun [18–20] and, more recently, the quiescent Sun [21, 22] during its nine-year lifetime. Having been constructed with point sources in mind, it boasts a sub-arcminute angular resolution and energy range of 3–78 keV.

In this Letter, we use a subset of NuSTAR’s recent quiescent solar limb dwells to constrain the axion-photon and axion-electron coupling, leveraging the characteristic spatial and spectral features of the basin signal. We compute the full production rate for a variety of production processes, and set limits using three different methodologies in order to adopt the most robust bound on parameter space. Two of these methods (Poisson and CL_s limits) are familiar, but we also present the first multidimensional implementation of a background-agnostic limit-setting methodology based upon the work of Yellin [23, 24]. We place constraints stronger than existing limits [25–28], by about an order of magnitude, on the axion-photon and axion-electron couplings, for axion masses between 5 keV and 30 keV. Figures 1 and 2 show our main results, taken along slices of parameter space given by $g_{a\gamma\gamma} = \alpha/2\pi f$ and $g_{aee} = m_e/f$ in Fig. 1 where f is the axion decay constant, and $g_{aee} = 0$ in Fig. 2.

The data and code used to obtain the results of this study is available on GitHub , and a link  below each figure provides the code with which it was generated. We employ natural units wherein $\hbar = c = k_B = 1$.

II. SIGNAL

We consider the leading axion couplings to electrons and photons:

$$\mathcal{L} \subset \frac{1}{2} \partial_\mu a \partial^\mu a - \frac{1}{2} m^2 a^2 \quad (1)$$

$$+ \frac{g_{aee}}{2m_e} (\partial_\mu a) \bar{\psi}_e \gamma^\mu \gamma^5 \psi_e - \frac{1}{4} g_{a\gamma\gamma} a F_{\mu\nu} \tilde{F}^{\mu\nu}, \quad (2)$$

*Electronic address: wderocco@stanford.edu

†Electronic address: srw9487@nyu.edu

‡Electronic address: bwgref@srl.caltech.edu

§Electronic address: jhuang@perimeterinstitute.ca

¶Electronic address: kenvt@nyu.edu, kvantilburg@flatironinstitute.org




FIG. 1: New limits (90%CL) on axion parameter space, with couplings determined by the axion decay constant f as $g_{a\gamma\gamma} = \alpha/2\pi f$ and $g_{aee} = m_e/f$. All regions shaded in red are excluded by our analysis. The thick red curve shows the Poisson limit (Sec. III B 1), the regular red curve shows the optimum cuboid (Yellin) method (Sec. III B 2), and the thin red curve shows the CL_s limit, which is the most stringent in most of the parameter space (Sec. III B 3). This suggests that the modeled smooth X-ray background dominates over the background from microflares in the data used in this analysis. The gray shaded regions are excluded by cooling of horizontal branch stars [26], white dwarfs [29], and red giants [30]. Note that the weakening of the CL_s limit at low masses is due to the use of a 4 keV lower photon energy cutoff to avoid solar contamination in comparison to the 3 keV cutoffs for Poisson and Yellin. \square

where m is the axion mass, a is the axion field, g_{aee} is the axion coupling to electrons (contributing primarily to axion production), and $g_{a\gamma\gamma}$ is the axion coupling to photons (allowing decays, and also contributing to production).

The energy density injection rate into the basin at a radius R as a result of these processes has been computed in general in Ref. [10], and can be put into the form

$$\dot{\rho}_b(R) \simeq \frac{3}{16\pi} \frac{GM_\odot}{R^4} \int d^3R' \tilde{Q}(R') \sqrt{|\Phi(R')|}, \quad (3)$$

where M_\odot is the mass of the Sun and $\Phi(R')$ is the gravitational potential at the site of production R' , with the integral evaluated over the solar volume ($0 < R < R_\odot$). The function $\tilde{Q}(R')$ encapsulates the production rate of non-relativistic particles into bound orbits [10] and is the sum of three primary components: Compton production (\tilde{Q}_C), bremsstrahlung (\tilde{Q}_B), and Primakoff production (\tilde{Q}_P). Expressions for the first two of these components appear in Ref. [10], and we derive the Primakoff expression

$$\tilde{Q}_P = g_{a\gamma\gamma}^2 \frac{\sqrt{2}\alpha n_e T}{\pi^2} \frac{m^4}{m_e e^{m/T} - 1} \quad (4)$$




FIG. 2: Limits on the axion-photon coupling $g_{a\gamma\gamma}$ with $g_{aee} = 0$. See Fig. 1 for details. \square

in App. A, where n_e is the electron number density, T the temperature, and α the fine structure constant. The quantity $\dot{\rho}_b(R)$ scales quadratically with g_{aee} (for the \tilde{Q}_C and \tilde{Q}_B terms) and $g_{a\gamma\gamma}$ (for the \tilde{Q}_P term); a plot of its dependence on m can be found in App. B.

The axions in the basin decay to photons ($a \rightarrow 2\gamma$) either directly via their tree-level photon coupling or at loop level via their electron coupling, at a rate:

$$\begin{aligned} \Gamma_{\text{rad}} &= \frac{g_{a\gamma\gamma}^2 m^3}{64\pi} + \frac{g_{aee}^2 \alpha^2 m^7}{9216\pi^3 m_e^6} \\ &\approx 2.38 \times 10^{-1} \text{Gyr}^{-1} \left(\frac{g_{a\gamma\gamma}}{10^{-12} \text{GeV}^{-1}} \right)^2 \left(\frac{m}{10 \text{keV}} \right)^3 \\ &\quad + 5.02 \times 10^{-11} \text{Gyr}^{-1} \left(\frac{g_{aee}}{10^{-13}} \right)^2 \left(\frac{m}{10 \text{keV}} \right)^7. \end{aligned} \quad (5)$$

If the axion's lifetime Γ_{rad}^{-1} is shorter than the Sun's age $t_\odot = 4.6$ Gyr, then the system reaches a steady state in which the rate of basin energy density injection ($\dot{\rho}_b$, which we take to be time-independent) is balanced by the losses due to axion decay ($\rho_b \Gamma_{\text{rad}}$). In general, the present-day signal flux is proportional to $\rho_b \Gamma_{\text{rad}}|_{t=t_\odot} = \dot{\rho}_b (1 - \exp(-\Gamma_{\text{rad}} t_\odot))$. We are primarily interested in the region $R_\odot < R < 2R_\odot$, where the majority of solar basin axions reside, and where secular perturbations from the planets can be neglected [31]. In App. B, we argue re-absorption of basin particles can also be ignored for the parameter space considered.

Since axions trapped in the solar basin have low velocities ($v < v_{\text{esc}}$), the decay to two photons takes place effectively at rest in the frame of the observer, so each photon acquires an energy $E_\gamma \simeq m/2$. The spectral signature is therefore a line at X-ray energies.

The angular signature follows from the characteristic $1/R^4$ dependence of the basin injection rate, a universal feature of injection into a $1/R$ potential [10]. Integrating along a non-Sun-crossing line of sight yields a $1/\theta^3$

fall-off in observed flux with increasing θ , where θ is the angle from the center of the Sun. Additionally, there is a doubling discontinuity of the signal at the solar limb ($\theta_\odot \equiv \arcsin[R_\odot/\text{AU}]$), where decays obscured by the Sun for $\theta < \theta_\odot$ become visible when $\theta > \theta_\odot$. A plot of the resulting angular template can be found in App. B.

Explicitly, we compute the number flux of decay photons per unit solid angle $d\Omega$ at Earth by integrating along the line of sight distance z :

$$\begin{aligned} \frac{dN}{dt dA d\Omega} &= \int_0^{z_*} dz \frac{z^2}{4\pi z^2} \frac{\rho_b \Gamma_{\text{rad}}}{E_\gamma} \\ &= \frac{\dot{\rho}_b(R_\odot)}{2\pi m} R_\odot^4 [1 - \exp\{-\Gamma_{\text{rad}} t_\odot\}] \int_0^{z_*} dz \frac{1}{R^4}, \end{aligned} \quad (6)$$

with $z_* \simeq \text{AU}(\cos\theta - \sqrt{\sin^2\theta_\odot - \sin^2\theta})$ for $\theta < \theta_\odot$, and $z_* = \infty$ for $\theta > \theta_\odot$. We neglect the $\sim 10^{-3} R_\odot$ -thick region above the solar surface in which the optical depth drops abruptly to zero as it is negligible in comparison to the spatial resolution of the detector. Rewriting the radius in terms of the line-of-sight distance and angle, $R = \text{AU}\sqrt{\sin^2\theta + (\cos\theta - z/\text{AU})^2}$, we can extract the angular template function $T(\theta) \propto \int dz R^{-4}$, normalized such that $T(0) = 1$. The signal flux can then be put in the more convenient form

$$\begin{aligned} \frac{dN}{dt dA d\Omega} &\equiv S_0 T(\theta) = \frac{\dot{\rho}_b(R_\odot) R_\odot}{6\pi m} [1 - e^{-\Gamma_{\text{rad}} t_\odot}] T(\theta) \\ &\approx \frac{2.75 \times 10^{-7}}{\text{s cm}^2 \text{ arcsec}^2} \left[\frac{\dot{\rho}_b(R_\odot)/m}{10^{12} \text{ cm}^{-3} \text{ Gyr}^{-1}} \right] [1 - e^{-\Gamma_{\text{rad}} t_\odot}] T(\theta), \end{aligned} \quad (7)$$

with S_0 the signal flux per unit solid angle at $\theta = 0$. The total number of expected signal events μ is the expression of Eq. 7 integrated over the exposure (corrected by livetime), ancillary response functions (ARFs), and field of view.

III. DATA AND ANALYSIS

A. Data processing

The data used for our analysis were collected by NuSTAR on Sep 12, 2020 during a series of quiescent limb dwells. Our data set is taken from the dwell with the least contamination from localized flares (Orbit 2) and further restricted to that orbit's combination of camera head units (CHUs) with maximal livetime (CHU12). (It is necessary to restrict the data to a particular combination of CHUs in order to ensure the consistency of the spatial coordinates in our analysis; when pointing at the Sun, CHU4 is blinded, hence an uncertainty of 2 arcmin is introduced on the relative pointing direction between different CHU combinations [32].) Our selection comprises about 1500 s of observations, during which the solar center underwent a 1.26 arcmin shift through NuSTAR's field of view. All events in NuSTAR's calibrated

energy range (3–78 keV) were recorded. (A spatial plot of the data can be found in Appendix C 1.)

The photon collection efficiency of NuSTAR over the field of view is quantified in a discrete 13×13 partition of subfields for both detectors (A and B), each with angular extent of arcmin \times arcmin and separate ancillary response function (ARF) computed using the extended-source functionality of the `nuproducts` pipeline [32] under the assumption of a spatially uniform background. These ARFs are effective collection areas that take into account detector effects such as aperture stop obscuration, detector absorption, and vignetting [32] that vary over the field of view.

B. Limits

To set constraints, we choose three separate methodologies, each with their own associated benefits and challenges. The first is a simple Poisson limit, the second is a generalization of the optimum interval method of Refs. [23, 24] (to our knowledge, the first multi-dimensional implementation of the algorithm over binned data), and the third is a likelihood-based CL_s method.

1. Poisson limit

The premise behind the Poisson limit is simple: the expected signal counts should not exceed the *total* recorded counts at some level of confidence. We identify the signal region by integrating over the observation time and spatial coordinates, but restricting to a narrow window in energy $E \in [m/2 - 2\sigma_E, m/2 + 2\sigma_E]$, where $\sigma_E = 0.166$ keV is the spectral resolution. (See App. D). Denoting by N the total number of recorded counts in that energy interval, and by μ the corresponding expected number of counts in a given signal model specified by $(m, g_{a\gamma\gamma}, g_{aee}, \alpha_\odot, \delta_\odot)$, the Poisson likelihood is given by

$$P(N|m, g_{a\gamma\gamma}, g_{aee}, \alpha_\odot, \delta_\odot) = \frac{\mu^N e^{-\mu}}{N!}. \quad (8)$$

The 90% confidence limit μ_{lim} is set at the value of μ , which depends on $(m, g_{a\gamma\gamma}, g_{aee}, \alpha_\odot, \delta_\odot)$, such that the cumulative probability for $\mu > \mu_{\text{lim}}$ is less than 0.1. In practice, we compute a Poisson limit on S_0 (Eq. 7) for each $(m, \alpha_\odot, \delta_\odot)$, and conservatively report the least stringent limit among a set of 317 discrete, uniformly-spaced solar positions $(\alpha_\odot, \delta_\odot)$ within 6 arcmin around the fiducial solar position. This simple limit-setting procedure suffers from the loss of spatial and temporal information, and sets the standard for the methods of Secs. IIIB 2 and IIIB 3.

2. Yellin limit

The ‘‘Poisson’’ upper limit we obtained in Sec. III B 1 stems from a comparison between the expected signal events and the observed events integrated over the entire field of view and observation time. However, the signal has a known spatial, spectral, and temporal dependence, which one can exploit to mitigate some of the background contamination, even for a completely unknown distribution of the background, as proposed by Yellin in Refs. [23, 24]. The method is based on selecting regions within the observed range that contain exceptionally few events in comparison to the signal expectation, or equivalently, exceptionally large regions that contain a given number of events. By construction, this method is relatively insensitive to parts of the signal region with high background (e.g. a flare). Yellin’s method [23] was initially designed for searches using a signal model with a known one-dimensional distribution, for data that contain few events (e.g. [33]). It is relatively straightforward to find the largest interval which contains a certain number of events, from which one can then find the most improbable of these intervals, the ‘‘optimum interval,’’ in the data by comparing to Monte Carlo (MC) simulations generated with *only signal events*. Hence, a background-independent limit can be obtained.

In the case of our analysis, the data exists in a four-dimensional space spanned by $(\alpha, \delta) \equiv (\text{RA}, \text{DEC})$ coordinates in the field of view, energy, and time. A main background from the Sun comes from solar (micro)flares, which are localized features in both space and time. Therefore, we generalize the optimum interval method to a multi-dimensional space as outlined in Ref. [24], by identifying the ‘‘optimum cuboid’’ in the four spacetime dimensions. This dramatically increases the computational complexity, requiring efficient binning and downsampling strategies discussed in App. F that preserve the power of the optimum cuboid method, while handling the high dimensionality and statistics of our data set. We perform Monte Carlo simulations of pure signal to compute the cumulative distribution functions (CDFs) for the largest cuboid volumes enclosing n events for $n = 0, 1, 2, \dots, N_{\text{data}}$. These CDFs can then be combined to place a 90% CL limit on the signal flux parametrized by S_0 in Eq. 7. As for the Poisson limit, we repeat the procedure for 317 different solar positions, and report the least stringent limit for each value of m .

The main advantage of this optimum cuboid method is that it is completely independent of any background model, making it ideal for data sets with unknown backgrounds.

3. CL_s method

A (typically) more stringent upper limit can be set with the CL_s method [34], which requires a sufficiently flexible model capable of capturing the background while

not overfitting. We make no attempt at modeling the intermittent background from solar (micro-)flares. The dominant non-solar background in most of the energy range of interest arises from cosmic X-rays that enter the detector at a glancing angle, never having passed through the optical bench. Subleading contributions arise from solar lines and internal components of the telescope. We choose to model these three backgrounds using the spectral shapes measured in Ref. [35], but allow their respective normalizations to float; a plot of the spectra is shown in App. E 2. We treat these backgrounds as spatially uniform, unaffected by the ARFs, since they are not focused by the X-ray optics. Additionally, we raise the minimum energy cutoff for the data to 4 keV to avoid contamination from solar activity. See Appendix E 1 for more detail on the background modeling procedure.

The likelihood can be expressed as a product of Poisson likelihoods over spectral and spatial bins. We add a Gaussian prior around the fiducial solar position with standard deviation $\sigma_{\odot} = 2$ arcmin. (See Appendix E 2 for the full expression.) For each axion mass m , the model parameters are therefore the signal flux S_0 , the initial solar position, and the normalizations of the three background components. We perform Markov Chain Monte Carlo (MCMC) analyses at each m , and then apply the CL_s method [34] on the resulting marginalized posterior for S_0 to place our constraint.

IV. RESULTS

The resulting limits are depicted in Figs. 1, 2, and 3. We represent our results in the three-dimensional parameter space spanned by $m, g_{a\gamma\gamma}$, and g_{aee} in a few different ways. The first is to unify the two couplings in terms of an axion decay constant f via the relations $g_{a\gamma\gamma} = \alpha/2\pi f$ and $g_{aee} = m_e/f$, corresponding to an axion with a tree-level electron coupling, and one-loop anomaly coupling to photons, as in DFSZ-type models [36, 37]. In this case, the axion is produced mainly through the axion-electron coupling (g_{aee}), and decays primarily through the axion-photon coupling ($g_{a\gamma\gamma}$). This yields the limits displayed in Fig. 1, up to one order of magnitude stronger than existing stellar cooling constraints [25–30].

We plot in Fig. 2 the limits in m - $g_{a\gamma\gamma}$ parameter space, with zero axion-electron coupling, where the axion-photon coupling is responsible for both the production and the decay of the axion. These limits also (approximately) apply to the case where the axion-electron coupling is generated at one-loop order, as in KSVZ-type models [38, 39]. Our analysis places a limit on an axion coupling exclusively to photons, over an order of magnitude below existing constraints from horizontal branch star cooling [26].

Finally, we plot our results in g_{aee} - $g_{a\gamma\gamma}$ space in Fig. 3, with contours corresponding to different axion masses, encapsulating all information at discrete mass values, and showcasing the scaling of our constraints in various lim-

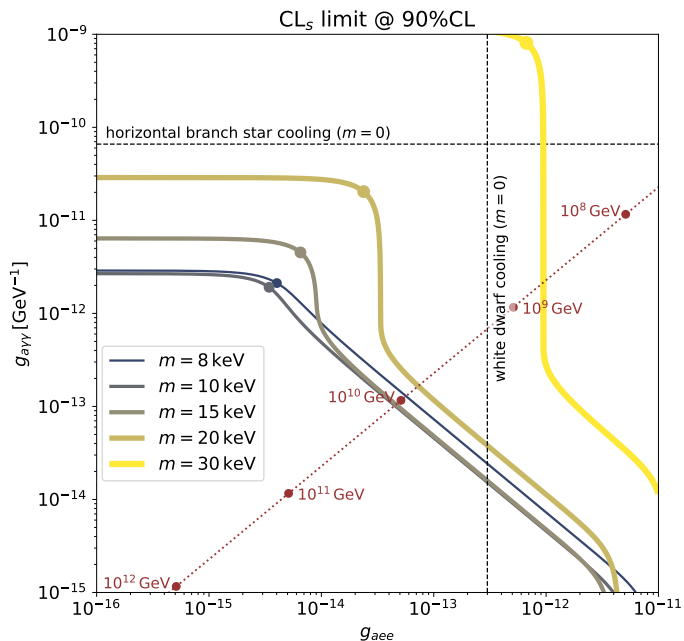


FIG. 3: The CL_s limits in the full three-dimensional parameter space, plotted on axes of $g_{a\gamma\gamma}$ and g_{aee} , with contours denoting different axion masses m . The red dotted line is the slice corresponding to the “universal coupling” displayed in Fig. 1. The horizontal and vertical dashed lines correspond to the cooling constraints (at $m = 0$) on the axion-photon coupling from horizontal branch stars [26] and on the axion-electron coupling from white dwarfs [29], respectively. \square

its. Tracing a single curve from the upper left to the lower right, one first sees the region in which the axion production is dominated by the Primakoff process, resulting in a bound independent of the electron coupling, and more stringent than horizontal branch star cooling constraints [26] for all but the highest mass displayed. At the dot, the Compton process overtakes as the dominant production mode, and our constraint is approximately independent of $g_{a\gamma\gamma}$. At even weaker photon couplings, the axion lifetime becomes longer than the age of the Sun, hence the basin decay has not yet fully reached detailed balance with the injection rate (see Eqs. 5 and 7), weakening the constraint to a contour at fixed product $g_{aee}^2 g_{a\gamma\gamma}^2$. Finally, at very small values of the photon coupling (lower right), the limits become vertical again, indicating the onset of axion decay via an electron loop, occurring well into the region constrained by white dwarf cooling [29] (vertical dashed line). The red dotted line corresponds to the “universal coupling” slice from Fig. 1.

The limits from the optimal cuboid (Yellin) and likelihood (CL_s) methods outperform the simple Poisson limit over the entire parameter space, as the former two are able to isolate the signal from the background. The optimum cuboid method offers the highest potential gain in data sets with localized, sporadic features such as flares. (In reality, the realized gain is artificially small, as we already picked a highly quiescent solar period.) Both the

Yellin and Poisson limits are completely independent of any background modeling. The likelihood limit typically outperforms the optimum cuboid limit since it includes background fitting of the temporally constant cosmic and instrumental X-ray backgrounds.

V. CONCLUSION

In this Letter, we have presented a first search for X-rays from axion decays in the Solar basin, using observations by the NuSTAR telescope. We set new constraints on axion parameter space by over an order of magnitude beyond existing bounds in the mass range of 5–30 keV.

Our analysis can be augmented by including a background model of the quiescent, micro-flaring Sun, which would strengthen the CL_s limit, and make possible a potential positive detection using the likelihood method. Longer, dedicated solar observations would also improve the sensitivity of both the likelihood and optimum cuboid methods. The optimum cuboid limit could also tighten with the inclusion of more (and possibly noisier) data, as it automatically selects the least likely region of observation space. Both analyses could also be marginally improved with a more precise knowledge of the fiducial position of the Solar center. A calibration of NuSTAR’s ARFs down to lower energies would also extend the sensitivity down to about 2 keV.

The Sun is not the only star that can produce an axion basin: compact remnants such as white dwarfs and neutron stars are attractive targets, as are stars with hot, dense cores, such as (super-)giants and horizontal-branch stars. For targets with extremely low backgrounds, the smaller signal fluxes from these more distant sources could be compensated by stacking multiple exposures. Such strategies could extend the first-of-its-kind indirect detection search for a stellar basin presented here to a larger mass range, even smaller couplings, and other weakly-coupled particles, such as dark photons [40] and milli-charged particles [41]. As we have shown in this Letter, indirect detection of stellar basins can probe heretofore unexplored parameter space for weakly-coupled particles. With the plethora of extensions outlined above, stellar basins will be an exciting target in the hunt for new physics beyond the Standard Model.

Acknowledgments

Acknowledgments—We thank Iain Hannah, Crystal Kim, and David Smith for discussion regarding NuSTAR data products and instrumental response. We are grateful to Robert Lasenby for sharing preliminary results upon which App. A is based, as well as many clarifying discussions and comments on the draft. We thank Asimina Arvanitaki and Matthew Johnson for comments on the draft, and Andrea Caputo, Joshua Foster,

Cristina Mondino, Benjamin Safdi, Edoardo Vitagliano, and Huacheng Yu for useful conversations. WD is supported by the U.S.A. Department of Energy, Grant No. DE-SC0010107. Research at Perimeter Institute is supported in part by the Government of Canada through the

Department of Innovation, Science and Economic Development Canada and by the Province of Ontario through the Ministry of Colleges and Universities. The Center for Computational Astrophysics at the Flatiron Institute is supported by the Simons Foundation.

-
- [1] J. N. Bahcall and R. K. Ulrich, *Rev. Mod. Phys.* **60**, 297 (1988), URL <https://link.aps.org/doi/10.1103/RevModPhys.60.297>.
- [2] L. Wolfenstein, *Phys. Rev. D* **17**, 2369 (1978), URL <https://link.aps.org/doi/10.1103/PhysRevD.17.2369>.
- [3] S. P. Mikheyev and A. Y. Smirnov, *Yadernaya Fizika* **42**, 1441 (1985).
- [4] G. G. Raffelt, *Stars as laboratories for fundamental physics: The astrophysics of neutrinos, axions, and other weakly interacting particles* (1996), ISBN 978-0-226-70272-8.
- [5] G. Raffelt and L. Stodolsky, *Phys. Rev. D* **37**, 1237 (1988).
- [6] G. G. Raffelt, *Lect. Notes Phys.* **741**, 51 (2008), [hep-ph/0611350](https://arxiv.org/abs/hep-ph/0611350).
- [7] H. An, M. Pospelov, and J. Pradler, *Phys. Lett. B* **725**, 190 (2013), 1302.3884.
- [8] J. Redondo and G. Raffelt, *JCAP* **08**, 034 (2013), 1305.2920.
- [9] E. Hardy and R. Lasenby, *JHEP* **02**, 033 (2017), 1611.05852.
- [10] K. Van Tilburg, arXiv preprint [arXiv:2006.12431](https://arxiv.org/abs/2006.12431) (2020).
- [11] S. Weinberg, *Phys.Rev.Lett.* **40**, 223 (1978).
- [12] F. Wilczek, *Phys.Rev.Lett.* **40**, 279 (1978).
- [13] R. Peccei and H. R. Quinn, *Phys.Rev.Lett.* **38**, 1440 (1977).
- [14] A. Arvanitaki, S. Dimopoulos, S. Dubovsky, N. Kaloper, and J. March-Russell, *Phys. Rev. D* **81**, 123530 (2010), 0905.4720.
- [15] P. Svrcek and E. Witten, *JHEP* **06**, 051 (2006), [hep-th/0605206](https://arxiv.org/abs/hep-th/0605206).
- [16] L. DiLella and K. Zioutas, *Astroparticle physics* **19**, 145 (2003).
- [17] F. A. Harrison et al. (NuSTAR), *Astrophys. J.* **770**, 103 (2013), 1301.7307.
- [18] K. Cooper, I. G. Hannah, B. W. Grefenstette, L. Glesener, S. Krucker, H. S. Hudson, S. M. White, D. M. Smith, and J. Duncan, *Monthly Notices of the Royal Astronomical Society* **507**, 3936 (2021), ISSN 0035-8711, <https://academic.oup.com/mnras/article-pdf/507/3/3936/40346090/stab2283.pdf>, URL <https://doi.org/10.1093/mnras/stab2283>.
- [19] L. Glesener, S. Krucker, I. G. Hannah, H. Hudson, B. W. Grefenstette, S. M. White, D. M. Smith, and A. J. Marsh, *The Astrophysical Journal* **845**, 122 (2017), URL <https://doi.org/10.3847/1538-4357/aa80e9>.
- [20] J. Duncan, L. Glesener, B. W. Grefenstette, J. Vievering, I. G. Hannah, D. M. Smith, S. Krucker, S. M. White, and H. Hudson, *The Astrophysical Journal* **908**, 29 (2021), URL <https://doi.org/10.3847/1538-4357/abca3d>.
- [21] A. J. Marsh, D. M. Smith, L. Glesener, I. G. Hannah, B. W. Grefenstette, A. Caspi, S. Krucker, H. S. Hudson, K. K. Madsen, S. M. White, et al., *The Astrophysical Journal* **849**, 131 (2017), URL <https://doi.org/10.3847/1538-4357/aa9122>.
- [22] S. Paterson, I. Hannah, B. Grefenstette, H. Hudson, S. Krucker, and L. Glesener, in *AAS/Solar Physics Division Meeting* (2020), vol. 52 of *AAS/Solar Physics Division Meeting*, p. 210.13.
- [23] S. Yellin, *Phys. Rev. D* **66**, 032005 (2002), URL <https://link.aps.org/doi/10.1103/PhysRevD.66.032005>.
- [24] S. Yellin, *Extending the optimum interval method* (2007), 0709.2701.
- [25] M. Tanabashi et al. (Particle Data Group), *Phys. Rev. D* **98**, 030001 (2018).
- [26] A. Ayala, I. Domínguez, M. Giannotti, A. Mirizzi, and O. Straniero, *Physical review letters* **113**, 191302 (2014).
- [27] F. Capozzi and G. Raffelt, *Phys. Rev. D* **102**, 083007 (2020), 2007.03694.
- [28] M. Giannotti, I. G. Irastorza, J. Redondo, A. Ringwald, and K. Saikawa, *JCAP* **10**, 010 (2017), 1708.02111.
- [29] M. M. Bertolami, B. E. Melendez, L. G. Althaus, and J. Isern, *Journal of Cosmology and Astroparticle Physics* **2014**, 069 (2014).
- [30] N. Viaux, M. Catelan, P. B. Stetson, G. Raffelt, J. Redondo, A. A. Valcarce, and A. Weiss, *Physical Review Letters* **111**, 231301 (2013).
- [31] C. Giovanetti, R. Lasenby, and K. Van Tilburg, in preparation (2022).
- [32] M. e. a. Perri, *The nustar data analysis software guide*, http://heasarc.gsfc.nasa.gov/docs/nustar/analysis/nustar_swguide.pdf.
- [33] S. Henderson, J. Monroe, and P. Fisher, *Phys. Rev. D* **78**, 015020 (2008), 0801.1624.
- [34] A. L. Read, *Journal of Physics G: Nuclear and Particle Physics* **28**, 2693 (2002), URL <https://doi.org/10.1088/0954-3899/28/10/313>.
- [35] D. R. Wik et al., *Astrophys. J.* **792**, 48 (2014), 1403.2722.
- [36] M. Dine, W. Fischler, and M. Srednicki, *Phys. Lett. B* **104**, 199 (1981).
- [37] A. R. Zhitnitsky, *Sov. J. Nucl. Phys.* **31**, 260 (1980).
- [38] J. E. Kim, *Phys. Rev. Lett.* **43**, 103 (1979).
- [39] M. A. Shifman, A. I. Vainshtein, and V. I. Zakharov, *Nucl. Phys. B* **166**, 493 (1980).
- [40] R. Lasenby and K. Van Tilburg, *Physical Review D* **104**, 023020 (2021).
- [41] A. Berlin and K. Schutz, arXiv preprint [arXiv:2111.01796](https://arxiv.org/abs/2111.01796) (2021).
- [42] G. G. Raffelt, *Phys. Rev. D* **33**, 897 (1986).
- [43] T. Altherr, *Annals Phys.* **207**, 374 (1991).
- [44] T. Altherr and U. Kraemmer, *Astropart. Phys.* **1**, 133 (1992).
- [45] T. Altherr, E. Petitgirard, and T. del Rio Gaztelurrutia, *Astropart. Phys.* **2**, 175 (1994), [hep-ph/9310304](https://arxiv.org/abs/hep-ph/9310304).
- [46] W. DeRocco, M. Galanis, and R. Lasenby (2022), 2201.05167.
- [47] K. J. H. Phillips, C. Chifor, and B. R. Dennis, *Astrophys.*

J. **647**, 1480 (2006), astro-ph/0607309.

- [48] B. W. Grefenstette et al., *Astrophys. J.* **826**, 20 (2016), 1605.09738.
 [49] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, *PASP* **125**, 306 (2013), 1202.3665.
 [50] L.-p. Ku, B. Liu, and W. Hsu (1997).
 [51] S. Nandy and B. Bhattacharya, *Computers & Mathematics with Applications* **36**, 11 (1998), ISSN 0898-1221, URL <https://www.sciencedirect.com/science/article/pii/S0898122198001254>.

Appendix A: Primakoff production rate

In this section, we compute the production rate of non-relativistic (NR) axions from the Sun through the axion-photon and axion-electron coupling. Axions can be produced via the Compton and bremsstrahlung processes, as well as the Primakoff and electro-Primakoff processes. The dominant production rate via the axion-electron coupling in the axion mass range of interest is the Compton process (see Ref. [10] for more details), with production function (to be used in Eq. 3):

$$\tilde{Q}_C = \frac{\alpha g_{aee}^2 n_e m^4 \sqrt{m^2 - \omega_{\text{pl}}^2}}{2^{3/2} \pi^2 m_e^4 e^{m/T} - 1}, \quad (\text{A1})$$

where $\omega_{\text{pl}} \equiv \sqrt{4\pi\alpha n_e/m_e}$ is the plasma frequency.

For the axion-photon coupling, the leading production mode is the Primakoff process depicted in Fig. 4, whose production rate vanishes at leading order in the NR expansion [42]. At this order, the only 3-momentum in the diagram is the momentum of the photon \mathbf{k}_γ and \mathbf{k}_a , so the matrix element of an $aF\tilde{F}$ coupling, depending on the cross product of \mathbf{k}_γ and \mathbf{k}_a , vanishes as $k_a \rightarrow 0$. At subleading order, there are two expansion parameters, the electron velocity $v_e \sim \sqrt{3T/m_e} \sim 0.1$ and the axion velocity $v_a \sim v_{\text{esc}} \sim 5 \times 10^{-3}$. In the following, we will compute the Primakoff production rate at $\mathcal{O}(v_e^2)$ in the NR expansion.

We start by expanding the spinors to first order in the velocity expansion. In the kinematic regime of interest, the incoming photon energy is approximately equal to the axion mass. Only the transverse mode of the photon contributes to the matrix element in vacuum, which is found to be

$$|\mathcal{M}|^2 = \frac{8g_{a\gamma\gamma}^2 e^2 m^2 (\mathbf{k}_i \cdot \mathbf{k}_\gamma)(\mathbf{k}_f \cdot \mathbf{k}_\gamma)}{\mathbf{k}_\gamma^4} \quad (\text{A2})$$

for $m_a \gg \max\{T, \omega_{\text{pl}}\}$, where \mathbf{k}_i and \mathbf{k}_f are the initial and final 3-momentum of the electron, and \mathbf{k}_γ is the 3-momentum of the incoming photon.

Extending the result to lower axion masses requires accounting for thermal effects [43–46]. The production

rate of an axion in the limit of vanishing v_a is

$$\Gamma = \frac{g_{a\gamma\gamma}^2}{16\pi^2 m} \int_{\omega_p}^{\infty} d\omega_\gamma k_\gamma \frac{\theta(q_0)}{e^{\omega_\gamma/T} - 1} \times \int_{-1}^1 d(\cos \phi) \beta(Q) g(K_\gamma, Q), \quad (\text{A3})$$

where $\beta(Q)$ is the imaginary part of the the propagator, $\theta(q_0)$ is a Heaviside theta function, and $g(K_\gamma, Q)$ is the coupling function:

$$g_L(K_\gamma, Q) = Q^2 k_\gamma^2 \sin^2 \phi \quad (\text{A4})$$

$$g_T(K_\gamma, Q) = (\omega_\gamma^2 q^2 + q_0^2 k_\gamma^2)(1 + \cos^2 \phi) - 4\omega_\gamma q q_0 k_\gamma \cos \phi$$

where g_L and g_T refer to the coupling functions of longitudinal and transverse plasmons, respectively. The ‘‘Cut’’ plasmon 4-momentum is $Q = (q_0, \mathbf{q})$, and that of the ‘‘Pole’’ plasmon $K_\gamma = (\omega_\gamma, \mathbf{k}_\gamma)$, while ϕ is the angle between the 3-momenta \mathbf{q} and \mathbf{k}_γ [43]. Following [43], we refer to the off-shell plasmon with momentum Q (see Fig. 4) as a ‘‘Cut’’ plasmon, while the on-shell initial state plasmon with momentum K_γ as a ‘‘Pole’’ plasmon. In the following, we will justify the approximations we made to arrive at Eq. A2 as well as its qualitative features by examining the properties of the functions $g(K_\gamma, Q)$ and $\beta(Q)$.

The emission rate depends on the coupling function $g(K_\gamma, Q)$, which contains the Lorentz contraction of the vertices in Fig. 4. The emission of relativistic axions via the Primakoff effect is dominated by the exchange of a longitudinal plasmon. However, the longitudinal coupling function g_L vanishes in the limit $\phi \rightarrow 0$ ($k_a \rightarrow 0$), signaling a v_a^2 suppression as expected (and consistent with the in-vacuum result). In contrast, the transverse coupling function g_T does not vanish in the same $\phi \rightarrow 0$ limit. In the following, we will therefore focus on the transverse-plasmon exchange.

The emission rate also depends on $\beta(Q)$, the imaginary part of the propagator of the ‘‘Cut’’ plasmon

$$\beta(Q) = -\frac{2\text{Im} \pi(Q)\theta(-Q^2)}{[Q^2 - \text{Re} \pi(Q)]^2 + [\text{Im} \pi(Q)]^2}, \quad (\text{A5})$$

where $\pi(Q)$ is part of the polarization tensor $\pi^{\mu\nu}(Q)$ as defined in Ref. [43] and computed in Ref. [46]:

$$\pi_L(Q) = \omega_{\text{pl}}^2 \frac{m_e}{q_0} \xi [Z(\xi_+) - Z(\xi_-)], \quad (\text{A6})$$

$$\pi_T(Q) = \omega_{\text{pl}}^2 \left\{ 1 - \left[\sigma^2 - \frac{Q^2}{4m_e^2} \right] \frac{m_e}{q_0} \xi [Z(\xi_+) - Z(\xi_-)] \right\},$$

with L and T subscripts referring to longitudinal and transverse plasmons. These results hold for a dilute NR plasma; see Ref. [46] for an in-depth analysis. In the above, we have employed the definitions $\xi_\pm \equiv \xi(1 \pm \frac{Q^2}{2q_0 m_e})$, the plasma dispersion function

$$Z(\xi) = \frac{1}{\sqrt{\pi}} \int_{-\infty}^{\infty} dx \frac{e^{-x^2}}{x - \xi} = i\sqrt{\pi} e^{-\xi^2} \text{erfc}(-i\xi), \quad (\text{A7})$$




FIG. 4: Feynman diagram for the Primakoff process.

with $\xi \equiv q_0/(q\sigma)$, and σ is the electron velocity dispersion [46]. The plasma dispersion function $Z(\xi)$ contains the entire imaginary part of $\pi(Q)$. We thus find:

$$\frac{\text{Im } \pi_T(Q)}{\text{Im } \pi_L(Q)} \sim \sigma^2 \quad (\text{A8})$$

for $Q^2 \ll (m_e\sigma)^2$, implying that the emission of relativistic axions through exchange of a transverse plasmon is suppressed by v_e^2 compared to the emission through exchange of a longitudinal plasmon, confirming our prior result.

This σ^2 difference between the longitudinal and transverse plasmon also shows up in the real part of $\pi(Q)$. Whereas the function $\pi_L(Q) \rightarrow \omega_{pl}^2/\sigma^2 \sim k_D^2$ in the limit $q_0 \rightarrow 0$, the function $\pi_T(Q) \rightarrow 0$ in the same limit. Here, $k_D = \sqrt{e^2 n_e/T}$ is the Debye screening wave vector. This suggests that the exchange of a transverse plasmon does not experience the same Debye screening as the exchange of a longitudinal plasmon, also as expected.

We can now integrate over the distribution of the initial state electron and photon, as well as the final state phase space to find the production rate of non-relativistic axions as

$$\tilde{Q}_P = g_{a\gamma\gamma}^2 \frac{\sqrt{2}\alpha n_e T}{\pi^2} \frac{m^4}{m_e e^{m/T} - 1}, \quad (\text{A9})$$

valid for axion masses $0.3 \text{ keV} \lesssim \omega_{pl} \ll m_a < 2m_e\sigma \approx 50 \text{ keV}$. As m_a approaches ω_{pl} , the production rate of non-relativistic axion through the ‘‘Cut-Pole’’ diagram becomes exponentially small due to the exponential suppression in the plasma dispersion function at large ξ .

The ‘‘Pole-Pole’’ contribution to the NR axion flux can only arise from the decay of the transverse plasmon to a longitudinal plasmon at the same momentum, but this rate is suppressed by the axion velocity v_a^2 and will not dominate in the mass range of interest to this work. NR axions can also be produced through the electro-Primakoff effect (‘‘Cut-Cut’’). Similarly, the matrix element can also only get contributions at $\mathcal{O}(v_e^2)$, so its rate is always subdominant compared to the Primakoff rate in the Sun.

Appendix B: Basin Evolution and Saturation

The energy density of axions in the solar basin depends not just on their injection rate from the Compton, bremsstrahlung and Primakoff production processes, but also on their effective accumulation time. There are three processes that contribute to basin depletion, and therefore set the effective lifetime of an axion in the stellar basin.

The first process is axion decay, with rate computed in Eq. 5, yielding a radiative lifetime within a few orders of magnitude from the age of the Sun (t_\odot) for the parameter space of interest. In the main text (Eq. 7 in particular), we have taken this radiative decay as the primary depletion process.

The second process is gravitational ejection, from close encounters or orbital resonances. For our data set based on observations of the solar limb, the vast majority of X-ray photons originate from axions on orbits with aphelia between R_\odot and $2R_\odot$. For such tight orbits, gravitational ejection times are much longer than the Sun’s age [31].

The third process is re-absorption of basin particles in the Sun, which can be estimated as [10]:

$$\tau_{\text{abs}} \simeq \frac{1}{\dot{\rho}_b(R)} m^4 f_{\text{sat}} \frac{v_{\text{esc}}^3(R)}{6\pi^2} \quad (\text{B1})$$

where $f_{\text{sat}} = \langle \frac{1}{e^{m/T} - 1} \rangle$ whose average is weighted by the production function \tilde{Q} over the solar volume, as in [40, Eq. 21]. This time is also much longer than Sun’s age for parameter space that is not already excluded for the mass range of interest, as shown in the bottom panel of Fig. 5, which assumes $f_{\text{sat}} = \frac{1}{e^{m/T_\odot} - 1}$ with constant $T_\odot = 1.3 \text{ keV}$ for simplicity, and evaluates τ_{abs} at $R = R_\odot$.

Appendix C: Data

1. Observation and detector response

We use data recorded by NuSTAR on 12-September-2020 during a series of quiescent limb dwells. Our dataset is taken from the dwell with the least contamination from localized flares (Orbit 2) and further restricted to that orbit’s combination of camera head units (CHUs) with maximal livetime (CHU12). We choose to restrict the data to a particular combination of CHUs in order to ensure the consistency of the spatial coordinates in our analysis; Solar observations are performed in NuSTAR’s Mode 06 (Science Mode) in which CHU4 is blinded, hence an uncertainty of ~ 2 arcmin is introduced on the absolute pointing direction between different CHU combinations [32]. This uncertainty also requires our analyses to marginalize over the true solar position. The effective exposure for CHU12 was 1501.17 seconds for focal plane module (FPM) FPMA and 1481.86 seconds for FPMB,




FIG. 5: *Top panel:* The axion density injection rate at the Solar surface, shown for two representative choices of axion, one coupled only to electrons (blue) and one to only photons (orange), as a function of axion mass. *Bottom panel:* Effective absorption time $\tau_{\text{abs}}(R)$ at the solar surface ($R = R_\odot$). Even when saturating the stellar cooling bounds for both the $g_{a\gamma\gamma}$ and g_{aee} couplings, the effective absorption time scale is longer than the age of the Sun t_\odot , especially for the mass range of interest for the analysis in this work $m \in [5, 30]$ keV. \square




FIG. 6: Angular dependence of axion decay signal flux, with θ measured with respect to the center of the Sun. Note the doubling at the solar limb and characteristic $1/\theta^3$ dependence for larger angles, which follows from the $1/R^4$ density profile of the Solar basin.

during which the solar limb underwent a $1.26'$ shift of $(\Delta\text{RA}, \Delta\text{Dec}) = (0.020^\circ, -0.0084^\circ)$ through NuSTAR's field of view. See Table I for further details of the observation.

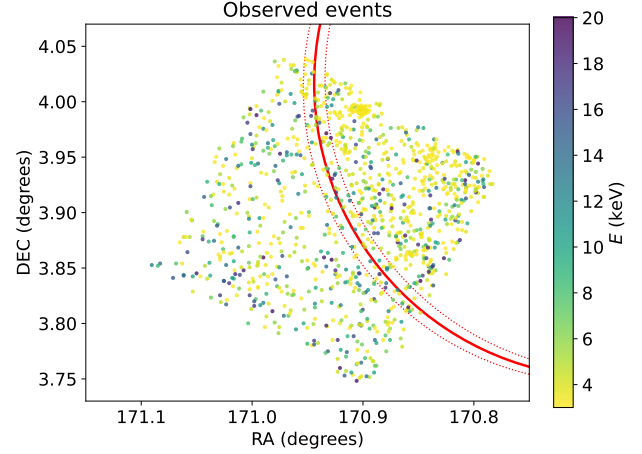


FIG. 7: Spatial map of the raw data used in our analysis, restricted to the relevant interval 3–20 keV. Each point represents a photon, and is colored according to its energy. The red solid line indicates the edge of the solar disk for the Sun at its fiducial position in the middle of the observation; the dotted red lines depict it at the start and end of the observations, with the Sun moving towards the bottom left. \square

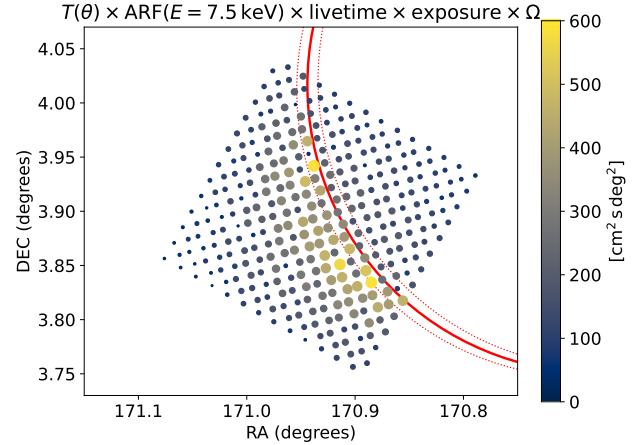


FIG. 8: Spatial map proportional to the expected number of counts in the 13×13 subfields of both detectors (for $m = 15$ keV). Specifically, we plot the product of: $T(\theta)$ with the Sun at its fiducial position in the middle of the observation window; the spatially varying ARFs at $E = 7.5$ keV; the effective exposure; the $\Omega = 1$ arcmin 2 solid angle of each subfield. \square

Event data was taken from Level 2 cleaned and calibrated event files. The spatial distribution is shown in Fig. 7. Energy was computed from the pulse invariant (PI) using the formula $E = (0.04 \text{ keV})\text{PI} + 1.6 \text{ keV}$. ARFs were computed using the `nuproducts` pipeline [32]. NuSTAR is primarily a point-source telescope, so the ARFs

Obs ID	Start (UTC)	Stop (UTC)	CALDB version	Exposure (s)
80610202001	2020-09-12T10:21:09	2020-09-12T11:56:09	20200912	1501.16599845754 (FPMA) 1481.86081041239 (FPMB)

TABLE I: Metadata of the NuSTAR observations used in this paper.

were computed using the extended source functionality and flat count distribution (assuming spatially uniform background) with no background extraction. Source regions were taken to be the cells associated with the gridding of the FOV into arcmin² regions (13×13 for both detectors) (see Sec. III A), with a different ARF calculated separately within each to account for the spatial dependence of the effective area due to detector effects (aperture stop, detector absorption, vignetting). The resulting expected spatial distribution of signal events is shown in Fig. 8.

2. Mock data generation

To compare to the true observation data, our limit-setting machinery was tested using mock data for which the axion mass and couplings are known. This mock data was generated in the following manner.

The background component was generated by sampling uniformly in time over the good time intervals for the observation and sampling energies from an interpolated distribution function formed from the 200-bin histogram of the true data. The background was sampled to be spatially uniform. The number of background counts was sampled according to Poisson statistics with expectation value given by the true number of counts in the data.

The signal component was generated by computing the expected number of counts of the signal model (Eq. 7) within each subgrid of the FOV for which an ARF was computed (taking the Sun to track its fiducial trajectory) and sampling from a Poisson distribution with that mean, then distributing those counts uniformly spatially over the subgrid and spectrally as a Gaussian with width σ_E centered at $E = m/2$. The temporal dependence was taken to be uniform. While some approximations made in the generation of this mock data result in slight differences from a true expected signal, the mock data sufficed to verify the efficacy of our modeling framework.

Appendix D: Energy Resolution

In order to determine a fiducial energy resolution for spectral lines in NuSTAR, we fit a known X-ray line in data of the active Sun with a Gaussian, taking the best fit width of the Gaussian σ_{line} as our fiducial resolution σ_E . We target the iron line at $E = 6.7$ keV [47] in observations of the active Sun taken on 01-November-2014 [48,





FIG. 9: Best fit curve to the 6.7 keV iron line in observations of the active Sun. The resulting best-fit width of the Gaussian component ($\sigma_E = 0.166$ keV) was taken as the fiducial NuSTAR energy resolution for the entirety of our analysis. 

Tab. 1]), and fit the spectrum with a function of the form

$$f(E) = c_1 \exp\left(-\frac{E}{c_2}\right) + c_{\text{line}} \exp\left(-\frac{(E - E_{\text{line}})^2}{2\sigma_{\text{line}}^2}\right) \quad (\text{D1})$$

where $\{c_1, c_2, c_{\text{line}}, \sigma_{\text{line}}\}$ were simultaneously fit to the data. The first term is included to model the solar background as a falling exponential, though since the line is narrow, the fit is insensitive to the particular functional form of this component. The second term models the line itself as a Gaussian centered at an energy E_{line} with width σ_{line} . The results of the fit are shown in Fig. 9. The fit yields a line energy $E_{\text{line}} = 6.45$ keV and a width $\sigma_{\text{line}} = 0.166$ keV, which we choose to use as our fiducial energy resolution σ_E for our analyses. This is a known marginal shift in the line energy (which is not yet understood), however the energy resolution is typical for NuSTAR observations [e.g., 17].

Appendix E: Likelihood Analysis Method

1. Background characterization

The background present in our data arises from several different sources, as mentioned in the main body of the text. The dominant background arises from cosmic X-rays that enter the detector at a glancing angle, never having passed through the optical bench, however sub-leading contributions arise from solar lines and internal components of the telescope. We choose to model the background using the spectral shapes measured in [35] but allowing the respective normalizations to vary. A plot of the various spectra is shown in Fig. 11. B_1 (“aper-

ture” in Fig. 9 of [35]) corresponds to the cosmic X-ray background that does not pass through the optical bench, B_2 (“int. lines”) to internal lines arising from the radiation environment of the detector, and B_3 (“int. cont.”) to gamma-ray Compton scattering and other featureless instrumental components. (See Appendix A of Ref. [35] for further discussion of these backgrounds.)

There is an additional contribution at very low energies that arises from active regions on the Sun. These stochastic features are not possible to encapsulate in a simple model, hence when performing the CL_s method, we have chosen to excise data contaminated by these microflares by raising the lower energy cutoff on the data used to set this limit. In order to select this cutoff, we performed a global fit of the background model to the observed data for E_{\min} ranging from 3 to 20 keV with a spacing of 0.24 keV. The χ^2 of the fit and coefficients B_1, B_2 , and B_3 stabilize near 3.8 keV, hence we selected 4 keV as our E_{\min} for the CL_s analysis. This cutoff explains the dramatic weakening of sensitivity of our CL_s bounds below $m \lesssim 8$ keV in Figs. 1 and 2. As can be seen in Fig. 12, with this choice of cutoff, the CL_s method is biased towards finding higher p -values, hence is a conservative methodology.

2. Likelihood computation

As stated in Sec. C 1, we grid the field of view into 13×13 grid of arcmin^2 cells, for both detectors A and B. The data, \mathbf{d}_{ijEt} , is therefore binned by spatial cell ($i, j = 1, \dots, 13$), energy, and time, with energy bin widths of 40 eV. Each of the time bins is multiplied by

the active exposure time of that bin. The data is also multiplied by the computed ancillary response function (ARF) that characterizes the detector effective area in each bin and the energy resolution $\sigma_E = 0.166$ keV. The Poisson likelihood for our flux model is a product over all bins:

$$P(\mathbf{d}_{ijEt}|m, g_{a\gamma\gamma}, g_{aee}, \alpha_\odot, \delta_\odot, B_1, B_2, B_3) \quad (\text{E1})$$

$$= \prod_{ijEt} \frac{\mu_{ijEt}^{n_{ijEt}} e^{-\mu_{ijEt}}}{n_{ijEt}!}$$

where μ_{ijEt} is the expected number of counts in each bin given the flux model, and n_{ijEt} is the observed counts.

The total flux model, including both the expected signal shape and backgrounds (denoted by f_k for functional form and B_k for free normalization, see Sec. E 1), is given by

$$\frac{dN}{dE dt} = B_1 f_{\text{aCXB}}(E) + B_2 f_{\text{internal}}(E) + B_3 f_{\text{cont}}(E)$$

$$+ \left[S_0 T(\theta) \frac{e^{-(E-m/2)^2/2\sigma_E^2}}{\sqrt{2\pi\sigma_E^2}} \right] \text{ARF} \times \Omega \quad (\text{E2})$$

where $\Omega = 1 \text{ arcmin}^2$ is the solid angle of each subfield. This leaves six free parameters: B_1, B_2 , and B_3 representing the background normalizations, S_0 for the overall signal strength, and α_\odot and δ_\odot for the (initial) solar position, which enters $T(\theta)$ through $\theta = \|(\alpha, \delta) - (\alpha_\odot(t), \delta_\odot(t))\|$. The angular flux template $T(\theta)$ is derived as shown in Sec. II and given explicitly by:

$$T(\theta) = \begin{cases} \frac{3\pi}{2} \sin^3 \theta_\odot \csc^3 \theta & \sin \theta_\odot < \sin \theta \\ \frac{3}{4} \sin^3 \theta_\odot \csc^2 \theta \left\{ \frac{4z_{\min}}{-1-2z_{\min}^2 + \cos 2\theta} + [\pi - 2\arctan(z_{\min} \csc \theta) \csc \theta] \right\} & \sin \theta \leq \sin \theta_\odot \end{cases}, \quad (\text{E3})$$

where $z_{\min} \equiv \sqrt{\sin^2 \theta_\odot - \sin^2 \theta}$. The expected signal is thus adjusted for the detector’s effective area per bin (via the ARF), and for the motion of the Sun during the data collection.

To set a limit we run a Markov Chain Monte Carlo (MCMC) which maximizes the likelihood function:

$$\sum_{ijEt} \text{LL}_m + \sum_{ijEt} \text{LL}_{\bar{m}} + \text{LP}_\odot \quad (\text{E4})$$

where LL_m is the logarithm of the likelihood given the data within $2\sigma_E$ of a given axion mass, and $\text{LL}_{\bar{m}}$ is that for the data outside of the $2\sigma_E$ range (separated such that only the *energy* spectrum of the non-axion region must be fitted to reduce computation time). We add a Gaussian prior to take into account our imprecise knowledge on

the solar position around its fiducial value ($\alpha_\odot^{\text{fid}}, \delta_\odot^{\text{fid}}$):

$$\text{LP}_\odot = -\frac{\cos(\delta_\odot^{\text{fid}})^2 (\alpha_\odot - \alpha_\odot^{\text{fid}})^2 + (\delta_\odot - \delta_\odot^{\text{fid}})^2}{2\sigma_\odot^2} \quad (\text{E5})$$

where $\sigma_\odot = 2 \text{ arcmin}$ is the uncertainty in the solar position. We also force B_1, B_2, B_3 and the data in any $ijEt$ bin to be strictly positive. We compute our MCMC chains via the `emcee` Python package [49], then extract a limit from the marginalized posterior on S_0 using the CL_s test statistic [34], defined by

$$\text{CL}_s = \frac{1 - p_{S_0}}{p_{S_0=0}} \quad (\text{E6})$$

where p_{S_0} denotes the cumulative distribution function at a given value of S_0 . A 90% CL upper limit on S_0 is set by solving $\text{CL}_s = 0.1$.





FIG. 10: Corner plot for MCMC fitting routine used in CL_s method at $m = 15$ keV. The units are arbitrary for the B_i , while S_0 is measured in units of $10^{-4} \text{ cm}^{-2} \text{ s}^{-1} \text{ deg}^{-2}$, and α_0 and δ_0 in degrees. Black (green) dashed lines in the 1D histograms correspond to 5%, 50%, and 95% quantiles for the posterior (prior). 

3. Diagnostic plots

Here we include various diagnostic plots for the CL_s fit for a particular value of axion mass, $m = 15$ keV. Fig. 11 is the ultimate spectral fit of the various background components. It can be seen that the best fit (blue) coincides well with the observed data (black). Fig. 10 shows the corner plot for the MCMC fitting procedure. Fig. 12 shows the ultimate confidence intervals on S_0 as a function of axion mass.

Appendix F: Optimum cuboid method

As discussed in Section III B 2, Yellin's original method [23] was designed for searches with a signal model with a known one-dimensional distribution, and for data containing very few events. This method was generalized [24] to searches with a larger number of events by searching for the largest interval (fraction of the data set) which contains no more than certain number of events. In this way, instead of computing the distance in the phase space between any two consecutive events, one can scan the interval size, and find the interval which sees a largest downward fluctuation in the number of events inside it. The method presented in Ref. [24] suggests the




FIG. 11: Best fit spectrum as determined during MCMC routine of CL_s limit at $m = 15$ keV. The associated background components are shown separately, with their sum shown in blue. B_1 corresponds to the aperture cosmic X-ray background, B_2 to internal lines, B_3 to internal continuum emission, and S_0 to the signal (see App. E 1). The true data are shown in black. \square

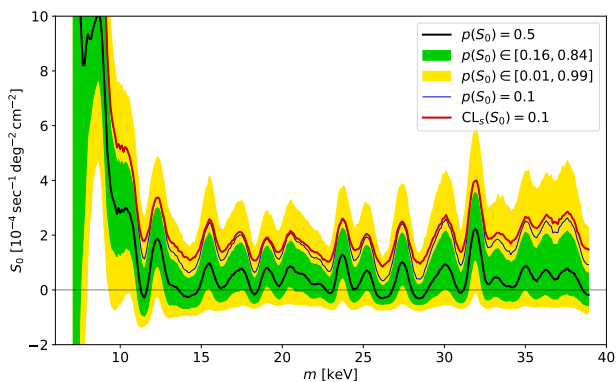


FIG. 12: Confidence intervals on S_0 using the CL_s method. The bias towards $S_0 > 0$ demonstrates that the model used is in general conservative. \square

possibility of coarsely discretizing the data region, which is essential in our generalization of the method to high statistics and to a higher-dimensional data space.

The main challenge for extending these methods to higher dimensional spaces is the rapid growth of computational complexity with the number of dimensions D . A naive expectation for the steps required for finding the largest cuboid with a certain number of data points n in a data set with N total points in D dimensions grows as $N^{2D-2}n$, while more advanced algorithms have reduced the complexity for finding the largest empty cuboid to $\mathcal{O}(N^{8/3}(\log N)^3)$ in 3 dimensions [50, 51]. For our data in $D = 4$, N ranges from $\mathcal{O}(10^2)$ at high axion masses to $\mathcal{O}(10^4)$ at low axion masses, where the complexity in the naive prescription is beyond our computational power.

As a result, we generalize the Yellin method to higher dimensions by discretizing the full experimental range into $k^4 = 10^4$ cuboids, constructed such that each cuboid contains exactly $1/k^4$ of the expected signal events. We then generate large number of MC samples for a fixed average of μ signal events (uniformly distributed on the unit cuboid). We count the number of signal events in

each unit cuboid and locate the largest cuboid $V_n^{(i)}$ which contains n number of events for the i th MC sample. This list $\{V_n\}^{(i)}$ over the many MC samples is then used to compute the cumulative distribution function of the V_n , denoted $C_n(V)$, for each n . In MC sample i , the maximum value of $C_n(V)$ achieves for any particular $V_n^{(i)}$ of that sample is denoted $C_{\max}^{(i)} = \max_n[C_n(V_n^{(i)})]$. A typical $C_n(V_n^{(i)})$ is around 0.5 by construction, as this corresponds to the median. A large $C_n(V_n^{(i)})$ indicates an upward fluctuation of a cuboid volume containing n points in sample i . Hence $C_{\max}^{(i)}$ is the least probable downward fluctuation in signal events in the sample. The list $C_{\max}^{(i)}$ can be used to compute \bar{C}_{\max} , the value for which the cumulative distribution function of $C_{\max}^{(i)}$ over all N_{sim} samples reaches 90%. The procedure is carried out for a list of μ values corresponding to signal strength with which we obtain a function $\bar{C}_{\max}(\mu)$, plotted in Fig. 13, which we can compare with data. (A glossary of these terms is provided at the end of this subsection for reference.)

We project each photon observed by NuSTAR, with coordinates (α, δ) , arrival time, and energy, onto a unit 4D cube. The projection is such that a *pure signal* distribution (i.e. Eq. 7) would be uniformly distributed on this unit 4D cube; we follow the projection implementation of Ref. [24, App. C]. We then subdivide this unit cube into $k^4 = 10^4$ bins as we did for the MC samples. For each axion mass m and each of the 317 solar positions samples, this procedure results in a list of events indexed by four integers indicating the generalized bin to which each event belongs. This list can then be used to compute a list of V_n^{data} , as well as $C_n(V_n^{\text{data}})$ of this list of V_n with the cumulative distribution functions obtained from the MC simulations for each μ . We can then find $C_{\max}^{\text{data}}(\mu)$ as the largest of $C_n(V_n^{\text{data}})$ of each μ value and compare it to $\bar{C}_{\max}^{\text{lim}}(\mu)$ of Fig. 13 for each μ value.

We have generated MC samples only for $\mu \leq 128$. If $C_{\max}^{\text{data}}(\mu) < \bar{C}_{\max}^{\text{lim}}(\mu)$ even for $\mu = 128$, we employ a




FIG. 13: The function (one minus) \overline{C}_{\max} as a function of the number of signal events as determined by Monte Carlo simulations (black points). Since \overline{C}_{\max} is computed from cumulative distribution functions, it has a statistical error which we estimate as $[N_{\text{MC}}(1 - \overline{C}_{\max}^{\text{MC}})]^{1/2}$, shown as 1- and 2-sigma bands in green and yellow. This error grows with increasing μ , but we compensate with a larger number N_{MC} of samples such that the fractional uncertainty in $1 - \overline{C}_{\max}$ is $\lesssim 10\%$ for all μ . For limit-setting purposes, we conservatively use the red line $\overline{C}_{\max}^{\text{lim}}$, which is its 2-sigma upward fluctuation. \square

down-sampling strategy. In that case, we randomly select 1/2 of the events projected on the 4D unit cube, and attempt to set a limit as outlined above on half of this data set. If $\mu = 128$ is still not excluded, we down-sample by another factor of 1/2, and so forth, in an iterative way. When this protocol converges after d down-sampling steps to a limit $\mu_{(d)}^{\text{lim}} < 128$, we take the limit on μ for the raw data to be $\mu^{\text{lim}} = 2^d \mu_{(d)}^{\text{lim}}$. Finally, we take the least stringent limit of the 317 solar positions at each axion mass. The same process is carried out for all axion masses to find the 90%CL limit μ^{lim} as a function of axion mass. This limit is plotted in the bottom panel of Fig. 15 as a red line, with a corresponding flux per unit solid angle at the solar center depicted in the top panel.

A list of the definition of the quantities we used to set limits are:

- k : Number of generalized bins per MC sample, in each dimension.
- $N^{(i)}$: Total number of simulated events in an MC sample, where i indexes the different samples.
- N_{MC} : Number of MC samples generated.
- $V_n^{(i)}$: The largest volume that contains no more than n points in the i -th MC sample.
- $\{V_n\}^i$: The array of $V_n^{(i)}$ for $n = 0, \dots, N^{(i)}$ in sample i . Note that $V_{N^{(i)}}^{(i)} = 1$ because the largest volume that contains no more than $N^{(i)}$ points is the entire box.
- $\mu = \langle N^{(i)} \rangle$: Expected number of data points in the MC samples. The following quantities will depend on μ implicitly.

- $P_n(V) = \text{PDF}_{V_n}(V)$: The probability distribution function, for a particular n , over the $V_n^{(i)}$ of the MC samples. The location of the peak of this PDF increases monotonically with n , since the corresponding average largest volume containing no more than n points necessarily gets larger.
- $C_n(V) = \text{CDF}_{V_n}(V)$: The cumulative distribution function of V_n , or simply $\int_0^V d\tilde{V} P_n(\tilde{V})$. The point at which this curve rapidly saturates towards 1 increases monotonically with n , see above.
- $C_{\max}^{(i)}$: In MC sample i , the maximum value that any $C_n(V)$ achieves for the particular $\{V_n\}^{(i)}$ of that sample: $C_{\max}^{(i)} = \max_n [C_n(V_n^{(i)})]$. Most $C_n(V_n^{(i)})$ would be expected to be around 0.5, as this corresponds to an average MC sample, hence a large $C_n(V_n^{(i)})$ indicates an upwards fluctuation in the large volume containing n points in sample i , $C_{\max}^{(i)}$ is the least likely such fluctuation across all n in the sample.
- $\Upsilon(C) = \text{PDF}_{C_{\max}}(C)$: The probability distribution function of C_{\max} over all N_{MC} samples.
- \overline{C}_{\max} : The value at which the CDF of $\Upsilon(C)$ reaches 0.9, i.e. $\int_0^{\overline{C}_{\max}} dC \Upsilon(C) = 0.9$.
- C_{\max}^{data} : This is the maximum of $C_n(V_n)$ across n on the observed data. The μ value for which $C_{\max}^{\text{data}} = \overline{C}_{\max}$ (note that both sides depend on μ) thus corresponds to our a 90%CL limit $\mu = \mu^{\text{lim, Yellin}}$.

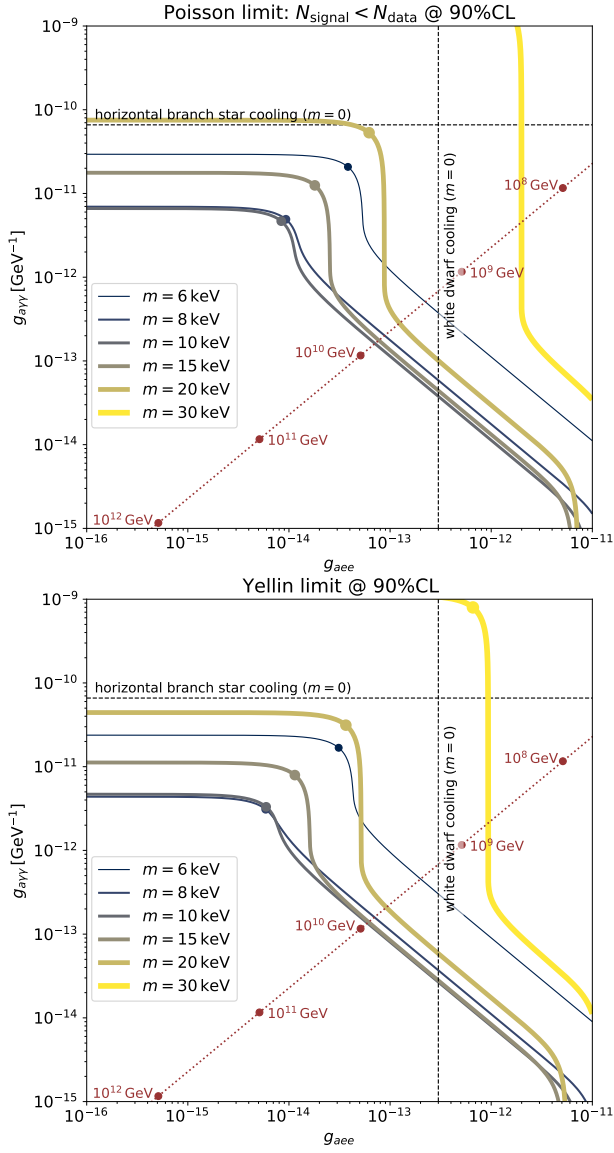


FIG. 14: The Poisson and Yellin limits in the full three-dimensional parameter space. Each yellow/brown curve corresponds to a particular axion mass, while black dashed lines indicate stellar cooling bounds on the couplings. The red dotted line is the slice corresponding to the “universal coupling” displayed in Fig. 1. See text of App. G for details. \square

Appendix G: Additional limit plots

In this section, we include a few additional constraint figures. Figure 15 is our “raw” exclusion figure upon which all others are based. It shows the limits on S_0 , the signal flux per unit solid angle at the center of the Sun, as a function of mass, for the three methodologies. The bottom panel shows the corresponding total number of events μ , in comparison to the total number N of events recorded.

In Fig. 14, we show the Poisson and Yellin limits in the full three-dimensional parameter space, the analogs




FIG. 15: *Top panel:* 90%CL limits on the signal flux per unit solid angle at $\theta = 0$, i.e. S_0 from Eq. 7, for the three methodologies (Poisson, Yellin, CL_s ; black, red, blue) used in this work. *Bottom panel:* Corresponding number of total expected photon counts μ for these three limit curves. Also plotted in gray is the total number of observed photons with $E \in [m/2 - 2\sigma_E, m/2 + 2\sigma_E]$. \square

of Fig. 3 for the CL_s limit. The limit curves have various turning points on this three-dimensional parameter space. As an example, consider the $m = 15$ keV line: going from left to right, the four distinct sections of the curve correspond to when

- The axion-photon coupling determines the production and decay rates of the axion (horizontal).
- The axion-electron coupling determines the axion production rate and the axion-photon coupling determines the axion decay rate, and the axion lifetime is shorter than the age of the universe (vertical).
- The axion-electron coupling determines the axion production rate while the axion-photon coupling determines the axion decay rate, and the axion lifetime is longer than the age of the universe (oblique).
- The axion-electron coupling determines the axion production and decay rates (decay proceeds via an electron loop), and the lifetime of the axion is longer than the age of the universe (vertical).

The second and fourth section may not show up for some axion masses.