



UvA-DARE (Digital Academic Repository)

Neutron stars as axion laboratories

Harnessing the power of the magnetosphere

Noordhuis, D.

Publication date

2024

[Link to publication](#)

Citation for published version (APA):

Noordhuis, D. (2024). *Neutron stars as axion laboratories: Harnessing the power of the magnetosphere*. [Thesis, fully internal, Universiteit van Amsterdam].

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Relativistic axions produced in neutron star polar-cap discharges

Axions can be copiously produced in localized regions of neutron star magnetospheres where the ambient plasma is unable to efficiently screen the induced electric field. As these axions stream away from the neutron star they can resonantly transition into photons, generating a large broadband contribution to the neutron star's intrinsic radio flux. In this chapter, we develop a comprehensive end-to-end framework to model this process from the initial production of axions to the final detection of radio photons, and derive constraints on the axion-photon coupling, $g_{a\gamma\gamma}$, using observations of 27 nearby pulsars. We study the modeling uncertainty in the sourced axion spectrum by comparing predictions from 2.5 dimensional particle-in-cell simulations with those derived using a semi-analytic model; these results show remarkable agreement, leading to constraints on the axion-photon coupling that typically differ by a factor of no more than ~ 2 . The limits presented here are the strongest to date for axion masses $m_a \simeq 10^{-8} - 10^{-5}$ eV, and crucially do not rely on the assumption that axions are dark matter.

This chapter is based on [261], and includes an updated discussion on the axion-photon conversion probability relative to the published work.

3.1 Introduction

In Chapter 2 we have investigated how dark matter axions can resonantly mix with photons in neutron star magnetospheres, where the large coherent magnetic field and the ambient plasma allow for this mixing to be highly efficient. As discussed, one expects this conversion to generate narrow spectral lines that can be observed using radio telescopes [197–212, 214, 219]; this idea has ignited numerous observational efforts [203, 206, 207, 209], with one of the most recent studies setting world-leading limits for axion masses near $\sim 30 \mu\text{eV}$ [209]. Despite the success and future promise of these spectral line searches, however, they are limited by the assumption that axions contribute significantly to dark matter, that the dark matter is smoothly distributed throughout the galaxy, and that the galactic dark matter halo has a cuspy profile. Furthermore, such searches are confined to axion masses $m_a \simeq 10^{-7} - 10^{-4} \text{ eV}$, as the mass must be sufficiently high to produce observable radio emission and sufficiently low so that resonances may be encountered.

Recently, Reference [213] proposed an alternative way to detect axions in neutron star magnetospheres that overcomes the aforementioned challenges. As mentioned in Chapter 1, this idea is based on axions sourced locally in the magnetosphere from the spacetime oscillations of $\vec{E} \cdot \vec{B}^1$, where \vec{E} and \vec{B} are the electric and magnetic fields of the neutron star. Reference [213] showed that the electromagnetic fields are strong enough in the polar caps of neutron stars (situated above the magnetic poles, and spanning distances of order $r_{\text{pc}} \sim \mathcal{O}(10 - 100)$ meters – see Figure 3.1) to produce an enormous flux of axions. As these axions traverse the magnetosphere they may encounter resonances similar to dark matter axions, generating a broadband radio flux in the MHz – GHz regime. This axion-induced radio flux provides an alternative observable in the search for these evasive particles.

In this chapter, we construct the first pipeline to compute the intrinsic spectrum of axions produced in neutron star polar caps, their resonant conversion to photons, and the non-linear evolution of these radio photons as they escape the magnetosphere. Our analysis uses both state-of-the-art numerical simulations as well as a newly developed semi-analytic model to predict the axion production rate; the overall agreement of these approaches illustrates that our procedure is not strongly sensitive to reasonable modeling uncertainties in the gap dynamics. We use this pipeline to constrain the axion-photon coupling by comparing the predicted radio flux with measurements of 27 nearby pulsars. The constraints derived here are the strongest to date for axion masses spanning $m_a \simeq 10^{-8} - 10^{-5} \text{ eV}$.

¹Axions may also be produced at the neutron star rotational frequency if there is a large-scale unscreened $\vec{E} \cdot \vec{B}$, see [262].

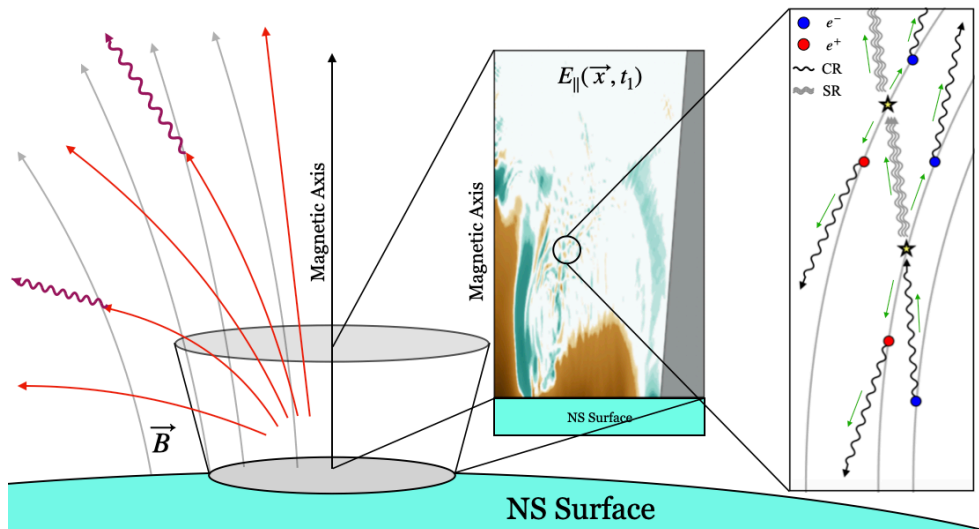


Figure 3.1: Schematic figure showing axion production in neutron star vacuum gaps. The vacuum gap is depicted by a truncated cone on the neutron star surface. Axions (red) are emitted from the gap and convert to photons (purple) in the presence of the neutron star’s magnetic field, \vec{B}_{NS} (gray). The left inset shows a time snapshot of E_{\parallel} (from the simulations of [191]), with the brown/green coloring reflecting negative/positive values of E_{\parallel} . The right inset depicts the microphysical processes responsible for the pair cascade, with green arrows indicating the direction the cascade flows with time.

3.2 Axion production from vacuum gap discharges

The e^{\pm} pair plasma populating the magnetospheres of neutron stars is expected to efficiently screen the component of the background electric field along the magnetic field lines (E_{\parallel}), except in small localized regions referred to as vacuum gaps which are responsible for particle acceleration and production of the pair plasma itself [184, 263]. Vacuum gaps are expected to arise in a variety of locations, including in the polar caps [187], the slot gap (located near the neutron star along the last closed field lines) [264], and the outer gap (located near the light cylinder along the last closed field lines) [265]. Recent progress in global particle-in-cell (PIC) simulations of the magnetosphere has shown that e^{\pm} pairs can be produced in the current sheet near the light cylinder, and efficiently screen gaps situated along the last closed field lines [231, 266]. We therefore choose to focus on the dynamics of the vacuum gaps in the polar caps, which also have the highest local values of $\vec{E} \cdot \vec{B}$.

Recently, local PIC simulations of the neutron star polar-cap pair cascade have been performed [188–191]. These simulations show that the cascade is an oscillatory process: e^{\pm} pairs accelerated in the gap produce gamma-rays that convert to more e^{\pm}

pairs, which proceed to screen the gap, shutting down pair production [267]. It was proposed that this process can produce coherent electromagnetic radiation, potentially answering the long-standing puzzle of the origin of pulsar radio emission [189, 192].

The oscillatory pair discharge process will in general lead to inductively driven oscillations of E_{\parallel} at a frequency set by the local plasma frequency, which is expected to evolve from the Goldreich-Julian (GJ) oscillation frequency $\omega_{\text{GJ}} = \sqrt{4\pi\alpha n_{\text{GJ}}/m_e}$ to a value $\mathcal{O}(10 - 100)$ times greater as the density of the pair plasma increases [190, 267–269]. Here we have introduced the GJ number density, given by $n_{\text{GJ}} = 2\vec{\Omega}_{\text{NS}} \cdot \vec{B}_s/e$ at the surface of a slowly rotating neutron star, with Ω_{NS} and B_s being the rotational frequency and surface magnetic field strength of the neutron star, e the elementary charge, m_e the electron mass, and α the fine structure constant. The presence of an oscillating E_{\parallel} directly enters as a source term in the axion’s equation of motion

$$(\partial_t^2 - \nabla^2 + m_a^2)a(x) = g_{a\gamma\gamma}(\vec{E} \cdot \vec{B})(x), \quad (3.1)$$

where a is the axion field and $g_{a\gamma\gamma}$ is the axion-photon coupling constant. Recall that this is a consequence of the generic coupling between axions and electromagnetism, captured via the Lagrangian term $\mathcal{L} = g_{a\gamma\gamma} \vec{E} \cdot \vec{B}a$.

To calculate an average axion production rate from the above, we follow the relevant discussion in [270]. We assume that the source is non-vanishing only over a time period T , which in our scenario corresponds to the quasiperiodic timescale of the gap collapse. In the asymptotic past the axion field can then be expanded as

$$a_0(x) = \int \frac{d^3k_a}{(2\pi)^3 \sqrt{2\omega_a(\vec{k}_a)}} \left(a_{\vec{k}_a} e^{-ik_a \cdot x} + a_{\vec{k}_a}^\dagger e^{ik_a \cdot x} \right), \quad (3.2)$$

where $a_{\vec{k}_a}^\dagger$ ($a_{\vec{k}_a}$) is the creation (annihilation) operator for the axion and $\omega_a(\vec{k}_a) = \sqrt{k_a^2 + m_a^2}$ is the axion energy. This can be thought of as the homogeneous solution to Equation 3.1 or, equivalently, the so-called ‘in’ state. The effect of the source, after it has been turned off, is to shift the creation and annihilation operators

$$a_{\vec{k}_a} \rightarrow a_{\vec{k}_a} + \frac{i}{\sqrt{2\omega_a(\vec{k}_a)}} \tilde{\mathcal{S}}(\vec{k}_a), \quad a_{\vec{k}_a}^\dagger \rightarrow a_{\vec{k}_a}^\dagger - \frac{i}{\sqrt{2\omega_a(\vec{k}_a)}} \tilde{\mathcal{S}}(-\vec{k}_a), \quad (3.3)$$

where $\tilde{\mathcal{S}}(\vec{k}_a)$ is the Fourier transform of the source term

$$\tilde{\mathcal{S}}(\vec{k}_a) = g_{a\gamma\gamma} \int d^4x e^{ik_a \cdot x} (\vec{E} \cdot \vec{B})(x). \quad (3.4)$$

The temporal integral here is performed over the period of gap collapse, while the spatial integral is performed over the gap volume. Equation 3.3 allows us to express

the average differential production rate of axions per momenta \vec{k}_a as

$$\frac{d\dot{N}_a}{d^3k_a} = \frac{\langle 0 | a_{\vec{k}_a}^\dagger a_{\vec{k}_a} | 0 \rangle}{(2\pi)^3 T} = \frac{|\tilde{\mathcal{S}}(\vec{k}_a)|^2}{2(2\pi)^3 \omega_a(\vec{k}_a) T}, \quad (3.5)$$

where $|0\rangle$ is the asymptotic vacuum state. The rate and spectrum of axions produced during the collapse of the vacuum gap is fully determined by the spacetime evolution of E_{\parallel} . The condition that axions be on shell inherently connects the spatial and temporal evolution of E_{\parallel} , meaning a given k_a -mode can only be produced if the spatial component of the Fourier transform contains support on scales $\sim k_a^{-1}$ and the temporal component contains support on scales $\sim \sqrt{k_a^2 + m_a^2}$.

In order to provide a quantitative understanding of axion production, we briefly estimate the production rate for a pulsar with $B_s = 10^{12}$ G and $\Omega_{\text{NS}} = 2\pi$ Hz. The polar-cap radius for such a pulsar is $r_{\text{pc}} \simeq 150$ m, and the maximum value of the unscreened electric field is roughly $E_{\parallel} \simeq 6 \times 10^{-6} B_s$ (see Appendix 3.A). We expect production to be most efficient when E_{\parallel} is largest, which occurs prior to the screening phase when the characteristic scale is $k_c \simeq 2\pi/r_{\text{pc}} \simeq 10^{-8}$ eV. Neglecting the phase in Equation 3.4, and taking $d^3k_a/\omega_a \simeq k_c^2$, $T = 10^{-7}$ s, and $g_{a\gamma\gamma} = 10^{-11}$ GeV $^{-1}$, we find $\dot{N}_a \simeq \mathcal{O}(10^{50})$ axions per second, which is comparable to the values obtained from more careful calculations in Appendix 3.A.

3.3 Modeling vacuum gap collapse

Here we discuss the two approaches taken to estimate the axion spectrum, Equation 3.5, produced by the periodic discharges of the vacuum gaps, one based on a semi-analytic model and the other on numerical simulations; a more detailed description of each step is deferred to Appendix 3.A.

3.3.1 Semi-analytic model

We begin by highlighting the key physical features entering our 2+1 dimensional semi-analytic model. We model two stages of gap evolution: pair cascade and gap collapse. The pair cascade phase begins with an initially unscreened E_{\parallel} and low charge density. As the pair cascade progresses, the number density increases exponentially until it reaches n_{CJ} . At this point the gap locally collapses, marking the end of the pair cascade.

A single seed particle (an electron) is assumed to initiate the pair cascade. Because of the presence of E_{\parallel} , the electron is accelerated to the radiation-reaction-limited Lorentz factor, γ_{max} , in time t_{acc} . Since particles move along curved field lines, they then emit curvature radiation (CR) photons with characteristic energy $\varepsilon_{\text{CR}} \propto \gamma_{\text{max}}^3/\rho_c$,

where ρ_c is the field line curvature [268, 271]. CR photons may be absorbed by the magnetic field to produce new pairs and synchrotron photons. Newly produced electrons (positrons) are accelerated away from (toward) the neutron star surface, producing their own CR photons, which may again produce new pairs. Synchrotron photons can also produce new pairs if their mean free path is less than the size of the gap (see the right inset of Figure 3.1 for an overview of the pair cascade).

In our model, we compute the creation positions, times, and energies of all photons and pairs produced by the single seed particle. Our model iterates over ‘generations’ of particles, where the first generation is the seed particle, and generation n particles are sourced by generation $n - 1$. We run five generations of the cascade for different seed particle locations within the gap, identifying points where the plasma density locally reaches n_{GJ} as the starting points for gap collapse – these points, which we refer to as ‘pair production seeds’, define the initial conditions of the semi-analytic model.

The initial stages of the gap collapse have been studied using analytic toy models and numerical simulations; in both cases, one expects the initial burst of particle production to induce exponentially damped oscillations of E_{\parallel} . These oscillations initially occur locally, and subsequently propagate outward along the magnetic axis [189–191, 233, 267, 269, 272–275]. The frequency of the damped oscillations is set by the local plasma frequency (which itself is set by the charge density and typical Lorentz factor), and will evolve from ω_{GJ} (*i.e.* the minimum value necessary to collapse the gap) to a value 10 – 100 times greater as more pairs are produced. The characteristic screening timescale τ_c is set by the time required for pair production processes to yield a GJ charge density, which is roughly equivalent to the time required to accelerate charges to γ_{max} [269]; for realistic pulsars this timescale lies between $10^{-9} - 10^{-6}$ s.

In order to capture the general features of the gap collapse process, we model the unscreened electric field E_{\parallel} with a static profile, and describe the screening of E_{\parallel} as the combination of outward propagating two-dimensional plane waves. Each plane wave is exponentially damped on a timescale τ_c , and has a time-evolving oscillation frequency growing from an initial value ω_{GJ} . In total, we source four propagating plane waves² evenly spaced across the width of the gap, and with initial conditions determined by the one-dimensional pair production process discussed above.

²This number must be sufficiently large to cover the gap, but cannot be too large as interference effects are not properly accounted for in this treatment. In Appendix 3.D we show that taking three or five plane waves leads to a negligible difference.

3.3.2 PIC simulation

Our second model used to compute the axion spectrum relies on a 2.5 dimensional³ PIC simulation developed in [190]. Working in axisymmetric cylindrical coordinates (r, z) , the authors impose a dipolar magnetic field. A rotating disk of radius r_{pc} is established at the stellar surface to produce a potential drop in the open field line zone of the neutron star. Outside $r = r_{\text{pc}}$, E_{\parallel} is forced to zero to model the plasma-filled closed field line zone. Particles are extracted from the surface at a rate that depends on the local value of E_{\parallel} ; these particles are accelerated to the radiation-reaction limit and emit gamma-rays through synchrocurvature radiation. Those gamma-rays subsequently produce e^{\pm} pairs in the ultrastrong magnetic field through one-photon magnetic pair production [276]. This process is modeled using the state-of-the-art quantum electrodynamics module in the PIC code OSIRIS [277]. Videos showing the dynamical screening of E_{\parallel} are available online [278]; a snapshot of the simulation is shown in the left inset of Figure 3.1.

Performing simulations from first principles of the gap collapse process is extremely challenging due to the large separation of scales between the size of the polar cap and the kinetic scale of the plasma (typically differing by 4 – 5 orders of magnitude). The simulations performed in [190] overcame this difficulty by re-scaling the quantum parameters $\chi_{\pm, \gamma} = \sqrt{(p_{\mu} F^{\mu\nu})^2} / (B_c m_e)$ governing both photon emission and pair production, multiplying them by a numerical constant; here, p_{μ} is the four-momentum of the corresponding e^{\pm} or γ , and $B_c \simeq 4.4 \times 10^{13}$ G is the Schwinger field strength. This re-scaling effectively allows pair production to occur at a much lower voltage drop. However, it also significantly reduces the inherent scale separation in the problem: the kinetic plasma length scale becomes only a few hundredths of the polar-cap size. As a result of this compression of scales, the Fourier transform necessarily compresses the power to an artificially narrower range of k_a -modes (the largest scales are expected to be unaffected, however power in small scales, *i.e.* large k_a -modes, will have been shifted to intermediate scales). We describe in Appendix 3.A a procedure for re-scaling the power of the Fourier transform in order to deal with this issue.

3.4 The radio flux

Once the initial axion spectrum has been determined, we employ an updated version of the ray-tracing algorithm discussed in Chapter 2 to compute the radio spectrum. We reiterate that one can calculate the radio flux generated via resonant axion-photon transitions through: (i) carefully following the trajectories of all axions created in the polar caps to the light cylinder, (ii) identifying the resonant conversion points, (iii)

³2.5 dimensional means using azimuthal symmetry to reduce the problem to two dimensions, but still evolving all three components of vector quantities.

computing the conversion probabilities from axions into photons, and (*iv*) propagating the produced photons through the dispersive magnetosphere to find their direction and energy at asymptotic infinity. The ray-tracing code from Chapter 2 was developed to treat the resonant conversion and propagation of radio flux arising from axion dark matter; we have modified this code here to include axion propagation in a Schwarzschild background and incorporated a generalization of the mixing which is valid for (ultra-)relativistic axions. Both of these new features are outlined below.

We calculate the rate of axion production for momenta spanning from the escape momentum to $\sim \mathcal{O}(10 \text{ GHz})$ (as mentioned in Chapter 1, the phenomenology of gravitationally bound axions differs markedly, and we leave a detailed study of bound states to Chapter 4). These momenta are used as initial conditions for the ray-tracing algorithm. Since axions are produced close to the neutron star, where gravity is strong, we propagate axions assuming they follow Schwarzschild geodesics⁴. The relevant equations describing the axion trajectories are

$$\frac{d^2 r}{d\tau^2} = -\frac{GM_{\text{NS}}}{r^2} + \frac{L^2}{m_a^2 r^3} - \frac{3GM_{\text{NS}}L^2}{m_a^2 r^4}, \quad (3.6)$$

$$\frac{d^2 \phi}{d\tau^2} = -\frac{2}{r} \frac{dr}{d\tau} \frac{d\phi}{d\tau}. \quad (3.7)$$

Here we have chosen to work in polar coordinates, with τ specifying the proper time, M_{NS} the neutron star mass, and G being the gravitational constant. It is important to note that both the angular momentum and total energy are constant along the trajectories and given by

$$L = m_a v_{\perp} \gamma r, \quad (3.8)$$

$$E = m_a \gamma \sqrt{1 - \frac{r_s}{r}}, \quad (3.9)$$

where γ is the standard Lorentz factor, v_{\perp} is the perpendicular component of the local particle velocity and r_s is the Schwarzschild radius of the neutron star.

The axions are propagated in just two dimensions, as their trajectories are inherently symmetric about the magnetic axis of the neutron star. The locations of the conversion points, however, are not. Consequently, we project each calculated axion trajectory in 70 different randomly sampled azimuthal directions when computing the conversion points and probabilities. This amount of samples was chosen to maximize stability and accuracy in the final results, while maintaining computational tractability. After

⁴This assumption does not hold inside of the neutron star itself. While our signal is not significantly affected by through-passing axions, we choose to re-scale the neutron star mass by the fraction of mass contained in a radius r (assuming a constant density profile) in order to avoid singularities. In the future one could improve upon this by adopting the interior Schwarzschild metric and taking the density profile under a fixed equation of state.

the ray-tracing, we average over the rotational period to compute the expected radio flux.

Resonant axion-photon transitions occur when $k_a \simeq k_\gamma$ [205], where the photon momentum is given by (this can be derived from the dispersion relation, Equation 2.1)

$$k_\gamma = \sqrt{\frac{\omega^2 - \omega_p^2}{1 - \frac{\omega_p^2}{\omega^2} \cos^2 \theta_{kB}}}. \quad (3.10)$$

Here, ω is the photon energy, ω_p is the plasma frequency, and θ_{kB} is the angle between the photon momentum and the external magnetic field \vec{B}_{NS} . In the non-relativistic limit, this resonance happens at $m_a \simeq \omega_p$; the generalized condition (which is also applicable in the relativistic regime), however, has the form

$$\omega_p^2 \simeq \frac{m_a^2 \omega^2}{m_a^2 \cos^2 \theta_{kB} + \omega^2 \sin^2 \theta_{kB}}. \quad (3.11)$$

We note that at the resonance, ω corresponds to both the axion and photon energy and $\theta_{kB} = \theta_B$ (with θ_B the angle between the axion momentum and the magnetic field).

In its most general form, the axion-photon transition probability is given by the Landau-Zener equation

$$P_{a \rightarrow \gamma} = 1 - e^{-\Gamma}. \quad (3.12)$$

This properly describes conversion in the adiabatic (*i.e.* large coupling and/or large magnetic field) regime, assuming that the length scale over which the mixing takes place is sufficiently small with respect to variations in the magnetic field and plasma density⁵. The adiabaticity parameter Γ corresponds to what we simply called $P_{a \rightarrow \gamma}$ in Chapter 2 (if one is not in the adiabatic regime that means $\Gamma \ll 1$ and thus $P_{a \rightarrow \gamma} \simeq \Gamma$), and we here use

$$\Gamma = \frac{\pi}{2} \left(1 + \frac{\omega_p^4 \xi^2}{\omega^4 \tan^2 \theta_B} \right) \frac{\omega^2 g_{a\gamma\gamma}^2 B_{\text{NS}}^2}{k_a} \frac{1}{|\omega_p \partial_s \omega_p + \frac{\omega^2 - \omega_p^2}{\omega^2 \tan \theta_B} \omega_p^2 \partial_s \theta_B|}, \quad (3.13)$$

where $\xi \equiv \sin^2 \theta_B / (1 - \omega_p^2 \cos^2 \theta_B / \omega^2)$ and $\partial_s \equiv \partial_{\parallel} - \omega_p^2 \xi / (\omega^2 \tan \theta_B) \partial_{\perp}$ (with ∂_{\parallel} and ∂_{\perp} respectively representing the derivatives in the parallel and perpendicular direction

⁵Recent work in [279] has shown that Equation 3.12 breaks down when the conversion is adiabatic and when these scales are comparable. For small axion masses, we do find that some photons are sourced from resonances in which Equation 3.12 does not hold; however, removing these photons entirely only leads to a modest correction to the constraints on the order of $\lesssim 10\%$ for axion mass $m_a = 10^{-9}$ eV and $\ll 1\%$ for $m_a \geq 10^{-8}$ eV, implying our constraints are robust to this uncertainty. We caution that lower mass axions are likely to be more severely affected by this problem, and thus a careful treatment would be required in this regime.

to the axion momentum at conversion). This formula is the relativistic generalization of Equation 2.10, and again all quantities should be evaluated at the resonance⁶. We addressed in Chapter 2 that the above conversion probability is not entirely correct; nevertheless, we continue to employ it here as the work discussed in this chapter was also published prior to this realization. We have again verified that this point, as well as the unnecessary inclusion of our dephasing treatment, should not alter our results beyond the already presented uncertainties. The derivation of the erroneous conversion probability, as well as the correct form, can be found in Chapter 5.

At every level-crossing we compute the conversion probability and axion survival probability⁷, $P_{a \rightarrow a} = 1 - P_{a \rightarrow \gamma}$, and subsequently weight each photon sourced at crossing i by $P_{a \rightarrow \gamma}^{(i)} = P_{a \rightarrow \gamma} \times \prod_{j=1}^{i-1} P_{a \rightarrow a}^{(j)}$. Photons are propagated to the light cylinder using the dispersion relation for a Langmuir-O mode (see Equation 2.1). After all photons have been traced, we determine the relative viewing angle of Earth, θ_{\oplus} , and compute the average radio flux by summing over the weighted contributions of photons with final positions $\theta \simeq \theta_{\oplus} \pm \epsilon$ (with ϵ taken to be very small and with photon frequencies falling in the desired bin $[f_1, f_2]$). This allows us to obtain final spectra of the axion-induced radio signal here on Earth, for any given neutron star. We provide an illustrate of the photon spectra produced by the pulsar B1737-30 (see Tables 3.1 and 3.2) in Figure 3.2, where the left panel shows a comparison between the predictions of the simulation and the semi-analytic model (at three viewing angles), while the right panel displays the observed flux density at five viewing angles (computed using only the semi-analytic model). In both cases, the axion mass is set to $m_a = 10^{-8}$ eV and the axion-photon coupling to $g_{a\gamma\gamma} = 10^{-11}$ GeV⁻¹. Figure 3.2 highlights two important points: there is reasonable agreement between the spectrum predicted by the PIC simulation and that of the semi-analytic model, and the observed flux density depends rather strongly on the observed viewing angle.

⁶Note that axion-photon conversion in the magnetospheres of highly magnetized neutron stars can become extremely efficient, generating order one conversion probabilities [209]. We pointed out in Chapter 2 that in the context of axion dark matter (in which axions are assumed to fall onto the neutron star), this can lead to an exponential suppression of the radio flux – this is a consequence of the fact that each axion is expected to pass an even number of resonances, and thus the generation of an observable radio signal requires a produced photon to survive a level crossing during escape from the magnetosphere (and $P_{\gamma \rightarrow \gamma} = 1 - P_{a \rightarrow \gamma} \ll 1$ when the conversion probability is large). We emphasize here that localized axion production in the polar caps naturally evades this suppression as these axions typically undergo an odd number of level crossings, making this a highly complementary probe of axions in the large coupling limit.

⁷We do not compute the axion production probability for subsequent resonances encountered by sourced photons, as this quickly becomes computationally prohibitive. For the pulsars and couplings studied here, we do not expect this secondary contribution to be significant.




Figure 3.2: *Left:* Comparison of photon spectra produced by pulsar B1737-30 using the PIC simulation (solid) and the semi-analytic model (dashed). Results are shown for three different viewing angles. *Right:* Photon spectra produced by pulsar B1737-30 using the semi-analytic model over a wider range of viewing angles. All results assume an axion mass $m_a = 10^{-8}$ eV and an axion photon coupling $g_{a\gamma\gamma} = 10^{-11}$ GeV $^{-1}$.

3.5 Results and conclusions

The radio emission mechanism of active pulsars is not well understood, making it difficult to identify signatures arising from this process. Nevertheless, one can constrain the existence of axions without knowing the intrinsic pulsar flux by requiring that their contribution not exceed the observed flux. In this section we derive limits on $g_{a\gamma\gamma}$ using observations of 27 nearby pulsars; our sample includes representative nearby pulsars, whose surface magnetic field and rotational period roughly span $B_s \simeq 10^{12} - 10^{13}$ G and $P_{\text{NS}} \simeq 0.1 - 2$ s. For computational ease, we choose to focus on pulsars whose radio emission geometry is constrained by observations, thus evading the need to marginalize over the misalignment and viewing angles. Details of all pulsars used in this analysis are presented in Appendix 3.B.

We present our fiducial 95% confidence level upper limits in Figure 3.3, which are obtained by computing the constraints for both models and taking the weaker limit at each mass (a comparison between models can be found in Figure 3.8). The bands around the fiducial limit represent a conservative estimate of the systematic uncertainties (see Appendix 3.D). For comparison, we plot constraints from radio line searches using neutron stars (blue), axion haloscopes (gray), the CAST experiment (teal), and X-ray and gamma-ray telescopes (light green). The limits derived in this chapter significantly improve upon existing bounds, and unlike axion haloscope experiments (and radio line searches), do not assume axions contribute to the dark matter. In addition, since the radio flux scales $\propto g_{a\gamma\gamma}^4$, the constraint is largely insensitive to minor mismodeling errors. The mass range covered by our constraints is limited by the frequency of radio observations (higher frequencies could probe higher




Figure 3.3: Upper limits on the axion-photon coupling derived in this chapter using a combination of a 2.5 dimensional PIC simulation and the semi-analytic model. The band reflects a conservative estimate of the modeling uncertainties (see Appendix 3.D). We compare to existing constraints from neutron stars [209] (blue), haloscopes [152, 160–167, 170, 280] (gray), helioscopes [172] (teal), and astrophysics [281–292] (light green). The former two have reduced opacity to highlight that they rely on axions being dark matter.

masses), and computational expense (computing time increases at both lower and higher masses).

In Appendix 3.C we show that the derived bound is controlled by observations of a few strong pulsars, with high frequency observations providing the most constraining power. A comprehensive analysis of all pulsars in the ATNF catalog, as well as more dedicated pulsar observations at high frequencies, could significantly improve upon these results; we reserve this broader analysis for future work.

Appendices

3.A Modeling vacuum gap collapse

In this section we illustrate specifics concerning the computation of the initial axion spectra. We begin by discussing in detail the relevant PIC simulation developed in [190], and then describe the ingredients used in the semi-analytic model.

3.A.1 PIC simulation

We use the results from one of the 2.5 dimensional quantum electrodynamics particle-in-cell simulations developed in [190] and performed using the code OSIRIS [277, 293]. Here, we outline the details of this simulation, introduce the re-scaling methodology, and explain how we extrapolate the simulation results back to more realistic regimes.

The simulation is performed in axisymmetric cylindrical coordinates (r, z) , where $r = 0$ is the rotation axis and $z = 0$ is the stellar surface (ignoring the curvature of the surface) modeled as a conductor. The stellar magnetic field is taken to be a dipole with magnetic axis aligned with the rotation axis, and magnitude given by $B(r = 0, z = 0) = 0.1 \times B_c$ with B_c the critical magnetic field strength (conventionally, $B_c \simeq 4.4 \times 10^{13}$ G, however in the simulation B_c is reduced to 10^7 G to make the computations feasible). To mimic the potential drop induced by the rotation of the neutron star, we rotate a disk of radius r_{pc} on the stellar surface; r_{pc} is the polar-cap radius and given by

$$r_{\text{pc}} \simeq R_{\text{NS}} \sqrt{R_{\text{NS}} \Omega_{\text{NS}}}, \quad (3.14)$$

where R_{NS} is the neutron star radius and Ω_{NS} its rotational frequency. The polar-cap radius is chosen to be $0.1 R_{\text{NS}}$ in the simulation, which corresponds to Ω_{NS} of a few hundred Hz. Rotation of the disk is achieved by imposing a boundary condition on the radial electric field $E_r(r, z = 0) = \tilde{E}_r \times (r/r_{\text{pc}}) \times g(r/r_{\text{pc}})$, with $\tilde{E}_r = \Omega_{\text{NS}} B_s R_{\text{NS}}$ and $g(x) = 0.5 \times (1 - \tanh((x - 1)/0.2))$. Here, B_s is the surface magnetic field strength of the neutron star. The latter function has been chosen to allow for a smooth transition from the open to the closed magnetic field line zone. The upper boundary of the

computational domain is taken to be open for both particles and fields [191, 277]. The plasma-filled closed field line zone is modeled as a perfect conductor where $E_{\parallel} = 0$.

The simulation is performed on a grid that spans $L_r \times L_z = 1.5 r_{\text{pc}} \times 2.5 r_{\text{pc}}$, with resolution 6000×10000 . This resolution, however, does not allow us to resolve the realistic scale separation between the polar-cap size r_{pc} and the plasma skin depth λ_p , especially when pair production creates a large number of e^{\pm} particles. In order to resolve the physics correctly, the plasma skin depth corresponding to the GJ charge density, $\lambda_{\text{GJ}} = \omega_{\text{GJ}}^{-1}$, is increased to a value of $\sim 5 \times 10^{-3} r_{\text{pc}}$, roughly $\mathcal{O}(100)$ times larger than expected in realistic neutron stars. This re-scaling decreases the overall potential drop induced in the open field line bundle. Nevertheless, in order to model the pair production physics the electrons must still be allowed to create pairs at the lower energies. This is achieved by also re-scaling the quantum parameters χ_{\pm} and χ_{γ} , which govern the emission of gamma-ray photons and the rate of conversion to e^{\pm} pairs. These quantum parameters are defined by $\chi_{\pm, \gamma} = \sqrt{(p_{\mu} F^{\mu\nu})^2} / (B_c m_e)$ – here p_{μ} is the four-momentum of the corresponding particle, $F^{\mu\nu}$ the electromagnetic field strength tensor, and m_e the electron mass – and are multiplied by a factor of $\zeta_{\pm, \gamma} = 10^3$ in the simulation to enhance pair production at lower energies.

The simulation proceeds as follows. At each time step in the simulation charges are injected at the neutron star surface with an effective charge density proportional to the local parallel electric field, $n_{\text{inj}} = \kappa(E_{\parallel}/er_{\text{pc}})$, with $\kappa = 0.1$ and e the unit of elementary charge. As the particles are accelerated from the surface, their quantum parameters χ_{\pm} are individually calculated, multiplied by the scaling factor ζ_{\pm} , and then used to emit synchrocurvature photons with energies sampled via Monte Carlo methods. Upon emission of a synchrocurvature photon, the energy of the photon is subtracted from that of the emitting particle. This is a discrete implementation of the ‘radiation-reaction force’ (see the following subsection for additional discussion). The emitted photons are also propagated in the simulation, their quantum parameters χ_{γ} are computed, multiplied by the scaling factor ζ_{γ} , and then used to compute the quantum Breit-Wheeler pair production cross section. The photon is afterward conditionally converted to an e^{\pm} pair based on the resulting pair production probability. The initial conditions of the simulation are set such that there are no particles along the open field lines. The conducting disk is gradually spun up, releasing particles which subsequently initiate quasiperiodic particle cascades. When performing the Fourier transform, we remove the initial part of the simulation, and focus only on the quasiperiodic solution achieved in the later stages.

The re-scaling of fundamental parameters mentioned above introduces a complication – the plasma scale is artificially inflated with respect to the size of the system (instead of being separated by $\sim 4 - 5$ orders of magnitude, they are only separated by ~ 3 orders of magnitude). When calculating the Fourier transform, this re-scaling shifts power from k_a -modes near $k_a \simeq \lambda_{\text{GJ}}^{-1}$ to a value $\mathcal{O}(100)$ times smaller. In order

to correct for this effect, we compute the Fourier transform \tilde{S} at a log-spaced series of momentum modes $(\tilde{k}_{\min}, \dots, \tilde{k}_{\max})$, and subsequently re-stretch the momentum vector $k_i = (\xi)^{i/n} \tilde{k}_i$ before computing the differential rate in Equation 3.5. Here, k_{\min} and k_{\max} are determined by the largest and smallest scales resolvable by the semi-analytic model, and n is the total length of the series. The extrapolation factor ξ is found by computing the ratio between the GJ scale predicted using the neutron star parameters and the value inferred from the ratio of $\lambda_{\text{GJ}}/r_{\text{pc}}$ used in the simulation.

Finally, note that, due to the parameter re-scalings highlighted above, one must re-scale the value of E_{\parallel}/B_s from the simulation to each pulsar under consideration. In order to do this, we use the fact that $E_{\parallel}/B_s \propto \Omega_{\text{NS}} r_{\text{pc}} \propto \Omega_{\text{NS}}^{3/2}$. In the simulation we analyzed, the parallel electric field, E_{\parallel}^s , corresponds to an effective rotational frequency of $\Omega_{\text{NS}}^{s,\text{eff}} \simeq 200$ Hz. Keeping the above in mind, we use the following re-scaling relation to determine the parallel electric field, E_{\parallel}^p , for any real pulsar

$$E_{\parallel}^p = \left(\frac{\Omega_{\text{NS}}^p}{\Omega_{\text{NS}}^{s,\text{eff}}} \right)^{3/2} \frac{B_s^p}{B_s^s} E_{\parallel}^s. \quad (3.15)$$

3.A.2 Semi-analytic model

There are two primary reasons to jointly develop a semi-analytic model. Firstly, the simulation discussed above was developed for a particular neutron star. We have applied it to the more general population by re-normalizing the polar-cap radius and the maximum amplitude of E_{\parallel} to their theoretical values, but it is not guaranteed that this re-normalization is sufficient to capture all relevant features. The second concern is that the re-scaling procedure introduced in the simulation to reduce scale separation may generate unexpected features in the Fourier transform. In order to address both of these concerns, we have built a semi-analytic model that can: (i) be applied to the broader class of pulsars and (ii) be run on a higher resolution grid, significantly reducing (although not entirely eliminating) the need for post-analysis stretching of the Fourier transform. We describe the details of this semi-analytic model below.

Pair cascade

An initially unscreened vacuum gap becomes unstable to runaway pair production. Due to variation in the magnetic field strength and line curvature across the gap, pair production occurs non-uniformly. Once the local charge density becomes comparable to the GJ value, dynamical screening of E_{\parallel} takes place, leading to exponential damping of the parallel electric field. The dynamics of E_{\parallel} in 1+1 dimensions are described in [269, 275] and summarized in the following section. In order to extend these results to 2+1 dimensions, we construct a 2+1 semi-analytic model of pair cascades in the gap.

The model provides the initial conditions for sourcing the aforementioned damped oscillations.

We use axisymmetric cylindrical coordinates with the lower boundary being a disk of radius r_{pc} rotating at frequency Ω_{NS} . The initial parallel electric field is constant throughout the gap with magnitude $E_{\parallel} = 2\Omega_{\text{NS}}B_s r_{\text{pc}}$. The domain of the gap is also permeated with a constant magnetic field $B_s \hat{z}$ ⁸. We further approximate the pair cascade initiated by a single seed particle to develop along a single magnetic field line, labeled by its radial coordinate r . The characteristic radius of a flux tube containing the particles produced in a cascade is of the order $\psi \sim 1/\varepsilon_{\gamma}$, where ψ is the angle between the photon momentum and the local magnetic field (*i.e.* the pitch angle) and ε_{γ} is the photon energy in units of m_e . Low energy photons that can have larger pitch angles generally have mean free paths greater than the size of the gap, preventing them from contributing to the cascade.

We inject macroparticles of charge κn_{GJ} at the neutron star surface, where $\kappa \ll 1$. Positrons that are accelerated away from the surface reach a maximum Lorentz factor determined by the balance between acceleration and radiative losses to curvature radiation according to the equation

$$m_e \dot{\gamma}(t) = \pm e \tilde{E}_r - \frac{2e^2}{3\rho_c(r)^2} \gamma^4(t), \quad (3.16)$$

where $+$ ($-$) applies to positrons (electrons), $\rho_c(r) = (4/3)R_{\text{NS}}^2/r$ is the magnetic dipole curvature at radius r , and $\gamma(t)$ is the particle Lorentz factor. The first term on the right-hand side corresponds to acceleration by the electric field and the second term to radiative losses due to emission of curvature photons. From Equation 3.16, the radiation-reaction-limited Lorentz factor is $\gamma_{\text{max}} = (3\tilde{E}_r \rho_c(r)^2 / 2e)^{1/4}$. A key quantity is the amount of time taken for a particle with initial energy γ_0 (in units of the electron mass) to reach energy γ_{max} . This acceleration time is largely independent of the initial Lorentz factor (see Figure 3.4). We define it as the amount of time it takes a particle to reach $\gamma = (1 - \epsilon)\gamma_{\text{max}}$, which can be determined by integrating Equation 3.16 to yield

$$t_{\text{acc}} = \frac{1}{4} \left(\frac{3m_e^4 \rho_c(r)^2}{2e^5 \tilde{E}_r^3} \right)^{1/4} \left(\pi + \ln \left(\frac{2}{\epsilon} \right) \right), \quad (3.17)$$

where $\epsilon \ll 1$. Going forward we set $\epsilon = 0.1$.

Once a particle reaches γ_{max} , it emits CR photons with characteristic energy $\varepsilon_{\text{CR}} = 3\gamma_{\text{max}}^3 / 2\rho_c$ and initial pitch angle $\psi_0 = 0$. As the photon propagates its pitch

⁸In reality the magnetic field strength varies across the gap, but the variation is of the order $r_{\text{pc}}/R_{\text{NS}} \ll 1$. The dominant contribution to the non-uniformity of pair production is the varying field line curvature.




Figure 3.4: Electron (solid) and positron (dashed) energies as a function of distance propagated along a curved magnetic field line with $\vec{E}_r = \Omega_{\text{NS}} B_s R_{\text{NS}}$. The neutron star parameters are taken to be $b \equiv B/B_c = 0.1$, $\Omega_{\text{NS}} = 2\pi$ Hz, and $\rho_c = 100$ km.

angle, and hence opacity, increases. The condition for absorption by the magnetic field is [268]

$$\tau(\varepsilon_\gamma, \ell_a) \equiv \int_0^{\ell_a} dx \alpha_B(\varepsilon_\gamma, \psi(x)) = 1, \quad (3.18)$$

where ℓ_a defines the mean free path and α_B is the opacity for pair production by a single photon in a strong magnetic field. Once the CR photon is absorbed it produces an e^\pm pair, each particle having energy $\varepsilon_\gamma/2$. Since $\psi_a \equiv \psi(\ell_a)$ is non-zero at absorption, the pair will be produced at a high Landau level and quickly decay to the ground state, emitting a number of synchrotron photons in the process. The final momentum (and energy assuming they are ultra-relativistic) of each fermion is given by the component of the parent CR photon energy parallel to the background magnetic field prior to absorption

$$\varepsilon_{\pm, f} = \frac{\varepsilon_\gamma}{2} \left(1 + \left(\frac{\chi_{\pm, a}}{b} \right)^2 \right)^{-1/2}, \quad (3.19)$$

with $\chi_{\pm, a} \equiv \chi_{\pm}(\ell_a)$ and $b \equiv B/B_c$. We assume the remaining energy is split equally among synchrotron photons of characteristic energy $\varepsilon_{\text{syn}} = 3b\psi_a\varepsilon_\gamma/8$. Once produced,

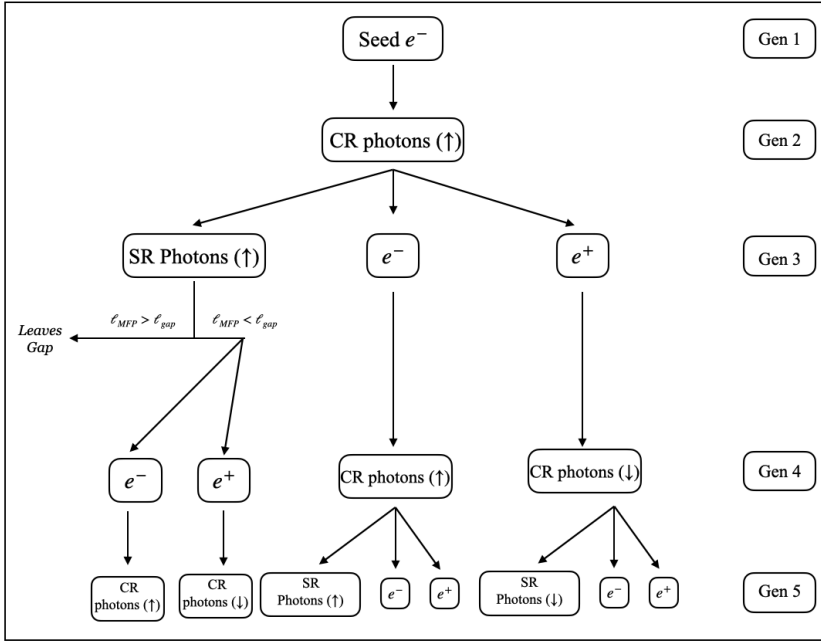


Figure 3.5: Schematic depiction of the generations of particle creation in our model. Photons are labeled as either upward propagating (\uparrow) or downward propagating (\downarrow). Each electron/positron produces many CR photons (see Equation 3.22 for an estimate). Each CR photon, when absorbed, produces a single electron/positron pair and several synchrotron photons.

the mean free path of synchrotron photons is calculated in the same way as for the CR photons, with the exception that they inherit the pitch angle of the parent CR photon. The secondary electron (positron) is emitted upward (downward) along the magnetic field line and reaches γ_{max} in time t_{acc} ⁹. The secondary electrons (positrons) can now emit upward (downward) propagating CR photons and the process continues.

We study the above model using an iterative algorithm that works as follows. We define four categories of particles: electrons, positrons, upward-propagating (γ^\uparrow) photons, and downward-propagating (γ^\downarrow) photons. For each e^\pm we keep track of its creation position (z_{create}^\pm) and time (t_{create}^\pm) and for each photon its creation position (z_{create}^γ), time (t_{create}^γ), energy (ε_γ), and initial pitch angle (ψ_0). We assume that particles reach γ_{max} at time $t_{create} + t_{acc}$, where t_{acc} is given in Equation 3.17, at which point acceleration is balanced by emission of CR. We assume CR photons are emitted at regular intervals $\delta t_{CR} = \varepsilon_{CR}/P_{CR}$, with $P_{CR} = 2e^2\gamma_{max}^4/3\rho_c^2$ being the CR power.

⁹The positron will take slightly longer to reach γ_{max} since it needs to turn around and re-accelerate in the opposite direction. For the secondary particles produced the turn-around time is small compared to t_{acc} (see dashed lines in Figure 3.4), so we neglect this complication.

Thus for each electron (positron) in generation n we create new upward-propagating (downward-propagating) CR photons according to

$$t_{\text{create}}^{\pm,(n+1)} = t_{\text{create}}^{\pm,(n)} + t_{\text{acc}} + i\delta t_{\text{CR}}, \quad (3.20)$$

$$z_{\text{create}}^{\pm,(n+1)} = z_{\text{create}}^{\pm,(n)} \pm t_{\text{acc}} \pm i\delta t_{\text{CR}}. \quad (3.21)$$

The characteristic number of CR photons emitted by each electron/positron is given by the transit time across the gap divided by the time interval between emission of CR photons

$$N_{\text{CR}} \simeq \frac{r_{\text{pc}} P_{\text{CR}}}{\varepsilon_{\text{CR}}} \simeq 400 \left(\frac{\gamma_{\text{max}}}{10^7} \right) \left(\frac{r_{\text{pc}}}{1 \text{ km}} \right) \left(\frac{1000 \text{ km}}{\rho_c} \right). \quad (3.22)$$

For each photon (indexed by i , with $i = 1, 2, \dots$) in generation n we compute the mean free path $\ell_i^{(n)}$ to pair production as described above and produce generation $n+1$ pairs and synchrotron photons at the point of absorption. We run five generations of particle production as described in Figure 3.5. The output of the algorithm gives information about the particle density as a function of position and time within the gap. We use this output to determine locations at which the charge density equals n_{GJ} , which serve as source points for gap collapse as described in the following subsection.

Screening of E_{\parallel}

The 2+1 dimensional semi-analytic model is constructed by fixing the spatial profile for the unscreened gap $E_{\parallel}(r, z)$, and mimicking the screening behavior induced by pair production processes using damped plane waves – notice that this is motivated by the qualitative features observed in the simulations of [191]. The physical scales driving the damping and oscillations of the plane waves are inferred using the simplified 1+1 dimensional analytic solutions, as well as the insight gained from the 2.5 dimensional PIC simulations. The procedure followed is outlined below.

We start by defining the physical dimensions of the vacuum gap. The radius of the polar cap r_{pc} is set by the typical footprint produced from a dipolar magnetic field and the height of the gap h_{gap} can be estimated by following the trajectories of e^{\pm} particles (which travel along magnetic field lines), and determining the point at which pair production generated a GJ charge density (as described above). This height is approximately given by [187]

$$h_{\text{gap}} \simeq 50 \text{ m} \left(\frac{\rho_c}{10^6 \text{ cm}} \right)^{2/7} \left(\frac{\Omega_{\text{NS}}}{1 \text{ Hz}} \right)^{-3/7} \left(\frac{B_s}{10^{12} \text{ G}} \right)^{-4/7}, \quad (3.23)$$

where ρ_c is the curvature of magnetic field lines near the neutron star surface. For a dipolar magnetic field, the radius of curvature is given by [294]

$$\rho_c \simeq 9.2 \times 10^7 \text{ cm} \left(\frac{\theta_{\text{pc}}}{\theta} \right) \sqrt{P_{\text{NS}}}, \quad (3.24)$$

where θ is the colatitude of the footprint of the magnetic field line, θ_{pc} the colatitude of the polar-cap boundary, and P_{NS} the rotational period of the neutron star. In determining the gap height we set θ to a constant characteristic value of $\theta_{\text{pc}}/2$ [294].

The size of E_{\parallel} may vary within the gap itself, growing linearly along the magnetic axis and decaying off at radial distances near r_{pc} . In order to incorporate smooth spatial boundaries we define a gap profile $\psi(r, z)$ such that the maximum unscreened electric field is given by $E_{\parallel, \text{max}}(r, z) = E_{\parallel, \text{max}} \times \psi(r, z)$, with $E_{\parallel, \text{max}} \equiv en_{\text{GJ}} r_{\text{pc}}$. We take the gap profile to be

$$\psi(r, z) = \frac{1}{2} \left[1 + \tanh \left(\frac{r - r_{\text{pc}}}{\delta r} \right) \right] \times \left(\frac{h_{\text{gap}} - z}{h_{\text{gap}}} \right), \quad (3.25)$$

with $\delta r = r_{\text{pc}}/10$. Notice that the radial dependence of ψ is somewhat ad hoc (although similar functional forms have been adopted *e.g.* in [191]). In reality, we simply require a quick and smooth fall-off of the gap near r_{pc} , and we expect other types of smooth functional cut-offs to provide comparable results. The linear growth of ψ on the other hand is intended to mimic the inability of the extracted surface current to screen E_{\parallel} , the efficiency of which should decrease as one moves away from the surface (see *e.g.* [188, 189]). We emphasize that the precise details of the gap profile do not have a large impact on the axion spectra.

The initial conditions of the semi-analytic model are that of a fully grown gap (*i.e.* $E_{\parallel} = E_{\parallel, \text{max}}(r, z)$). At this point, pair production is expected to collapse the gap, sending $E_{\parallel} \rightarrow 0$; the non-linear interplay between the motion of the plasma and the response of the electric field, however, induces damped oscillations in E_{\parallel} , with each phase of the oscillation triggering the production of higher generation e^{\pm} pairs. The damping is expected to be exponential, with a typical timescale roughly given by [269]

$$\tau_c \simeq 3.1 \times 10^7 \left(\frac{B_s}{10^{12} \text{ G}} \right) \left(\frac{10^5}{M} \right)^{1/2} \left(\frac{0.1 \text{ s}}{P_{\text{NS}}} \right)^{7/4} \frac{m_e}{eE_{\parallel, \text{max}}}, \quad (3.26)$$

where M is the charge multiplicity. The frequency of oscillation is set by the local plasma frequency. We assume here that the plasma frequency during the initial response is set by the GJ value (*i.e.* the minimal value required in order to suppress the electric field), and increases by a factor of 100 as subsequent generations of particles are produced (see *e.g.* [269]). The growth of the oscillation frequency should be slightly less than exponential, since only a subset of each generation's e^{\pm} pairs contributes in producing the subsequent generation (at this point, pair production is driven only

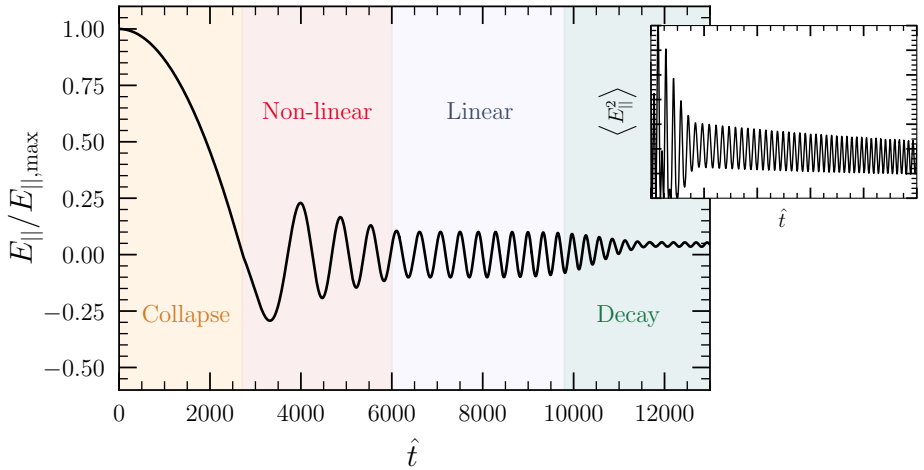


Figure 3.6: Illustrative depiction of the various phases (excluding the growth phase) of the evolution of E_{\parallel} in the vacuum gap, following the 1+1 dimensional formalism of [269]. The inset shows a zoom-in of $\langle E_{\parallel}^2 \rangle$ during the decay phase.

by the high energy tail of the distribution). As such, we adopt a time-dependent evolution of the oscillation frequency given by

$$\tilde{\omega}(t) = \omega_{\text{GJ}} \times (1 + \beta e^{(t/\tau_c)^\alpha}). \quad (3.27)$$

Here $\beta \equiv (\hat{\omega}_f - 1) \times e^{-2\alpha}$, and we set $\alpha = 0.5$ and $\hat{\omega}_f = 100$. In the following section we illustrate that this choice does not have a major impact on the derived constraints.

As an informative illustration, we show in Figure 3.6 the temporal evolution of the 1+1 dimensional solution identified in [269] (computed using the same parameters as used in their Figure 2) as a function of (re-normalized) time \hat{t} , and using the frequency evolution in Equation 3.27. In the figure, we extend the result beyond the non-linear damping regime and include the linear oscillation and decay regimes identified in [269]. We note that the latter two regimes are not significant for our purposes, as the solution at these late times is only expected to be representative of small, localized plasma overdensities located well outside of the polar cap.

The fundamental difficulty in modeling the gap collapse processes arises from the non-linear response that appears in two dimensions. Particle production itself is inherently inhomogeneous since pair production depends heavily on the radius of curvature of the magnetic field lines, which varies across the gap. The response of the electromagnetic fields to the inhomogeneity in particle production further induces non-linear dependencies, which significantly complicate the evolution of E_{\parallel} . Here, we attempt to capture the initial inhomogeneity in the pair formation front by

accelerating test particles and tracing the radiation and initial pair production events (see subsection on the pair cascade above); this procedure neglects the back-reaction on the electromagnetic fields, and thus is only valid when E_{\parallel} is large.

We model the subsequent non-linear screening of E_{\parallel} by drawing well-isolated samples from the pair formation front, and inducing two-dimensional damped plane waves that propagate outward – we refer to these samples as ‘pair production seeds’, with coordinates denoted by $(r_{\text{seed}}^i, z_{\text{seed}}^i)$. Note that similar behavior can be seen in the simulations of [190]. We approximate the evolution of E_{\parallel} using the Ansatz

$$E_{\parallel}(r, z, t) = E_{\parallel, \text{max}} \times \psi(r, z) \times \prod_{i=1}^{N_{\text{pts}}} \phi_{\text{pp}}^i(r, z, t), \quad (3.28)$$

where the product runs over each of the sampled seeds. Our intent with this model is to capture the typical temporal and spatial variations of E_{\parallel} . We will validate it below by comparison with PIC simulation results. The response of each seed is given by

$$\phi_{\text{pp}}^i(r, z, t) = \begin{cases} \cos(\tilde{\omega}\tilde{t}_i - k \cdot d_{\text{pp}}^i) e^{-\tilde{t}_i/\tau_c} & \text{if } \tilde{t}_i \geq 0, \\ 1 & \text{if } \tilde{t}_i < 0. \end{cases} \quad (3.29)$$

We take $k = \tilde{\omega}$ (*i.e.* assume it travels at the speed of light), and the functional dependencies on spacetime are understood to be implicit. We have introduced in Equation 3.29 the notion of an ‘effective time’ \tilde{t}_i which captures the relative spacetime dependent response of E_{\parallel} (in other words, this captures the relative delay that would arise in the event that gap collapse occurs locally and expands outward, as seen in [190]); this is given by

$$\tilde{t}_i(r, z, t) \equiv t - t_{\text{pp}}^i - d_{\text{pp}}^i(r, z, t), \quad (3.30)$$

where t is the absolute time, t_{pp}^i is the initial time at which particle production event i was initiated, and d_{pp}^i is the relative distance from the center of particle production event i to the position (r, z) . Since the electric field will push the newly produced plasma out of the gap, we adopt a drift in the initial seed location such that

$$z_{\text{pp}}^i = \begin{cases} z_{\text{seed}}^i & \text{if } (t - t_{\text{pp}}^i) < t_{\text{delay}}, \\ z_{\text{seed}}^i + (t - t_{\text{pp}}^i - t_{\text{delay}}) & \text{otherwise.} \end{cases} \quad (3.31)$$

The factor t_{delay} is set by the minimum time required for the plane waves to collapse the gap at all radii, which is also roughly the time over which the pair formation front forms, *i.e.* $\max(t_{\text{pp}}^i) - \min(t_{\text{pp}}^i)$ (note that the inclusion of t_{delay} is necessary to ensure that the gap collapses). We run each semi-analytic model for a total time set by $t_{\text{max}} = 2 \times \max(t_{\text{pp}}^i)$, which is sufficiently long to observe the gap collapse and




Figure 3.7: Number of axions produced per unit time predicted by the simulation (left panel) and semi-analytic model (right three panels) for various extrapolation factors ξ , ranging from 6 to 20. The spectra are produced for pulsar B1737-30 (see Tables 3.1 and 3.2), an axion mass of $m_a = 10^{-8}$ eV, and an axion-photon coupling of $g_{a\gamma\gamma} = 10^{-11}$ GeV $^{-1}$.

the subsequent outward drift of the seed points; moreover, this number agrees quite well with the expected gap collapse period of typical neutron stars, which is around $10^{-9} - 10^{-6}$ s [275].

In order to be able to perform the full Fourier transform across the momentum range of interest, one must have sufficient spatial and temporal resolution. Specifically, we require a spatial resolution at the level of the GJ oscillation wavelength (which is well beyond the resolution needed to resolve momentum modes $k_a \sim \mathcal{O}(10 \text{ GHz})$), and temporal resolution given by

$$\delta t = \min\left(\frac{2\pi}{\omega_{\text{GJ}}}, \frac{2\pi}{\omega_{\text{max}}}\right), \quad (3.32)$$

where ω_{max} is the maximum axion energy computed. For most of the pulsars and masses of interest, storing the gridded evolution of the gap at this resolution requires $\mathcal{O}(100 \text{ Gbs})$ to $\mathcal{O}(\text{Tbs})$ of memory.

In order to make these computations more feasible, we perform the semi-analytic calculations with a slightly enhanced value of the GJ scale, $\tilde{\lambda}_{\text{GJ}} = \xi \times \lambda_{\text{GJ}}$, with ξ taken to be 6 (roughly the smallest value which we can achieve given current resources). Similar to the case of the PIC simulation, this modification effectively inflates small scales and thus misplaces the power of the Fourier transform in smaller k_a -modes (here, we resolve around 4.5 orders of magnitude in scale, but require around 5 – thus the effect is rather minimal). As before, we attempt to correct for this by first computing $\tilde{\mathcal{S}}$ at a series of log-spaced momentum modes $(\tilde{k}_{\text{min}}, \dots, \tilde{k}_{\text{max}})$, and subsequently extrapolating the momentum vector $k_i = (\xi)^{i/n} \tilde{k}_i$ before calculating the differential rate in Equation 3.5. To illustrate the sensitivity to this extrapolation factor, we plot in Figure 3.7 the axion spectrum predicted using $\xi = 6, 10,$ and 20 for pulsar B1737-30 and an axion mass of 10^{-8} eV (the axion-photon coupling is set to

$g_{a\gamma\gamma} = 10^{-11} \text{ GeV}^{-1}$). In the left panel of Figure 3.7 we also show for comparison the corresponding spectrum produced by the simulation. One can see that the relative difference between the simulation and semi-analytic model greatly exceeds the small differences that appear as one changes the re-scaling factor. All runs are computed at 35 log-spaced k_a -values in both k_r and k_z . In Appendix 3.D, we return to the effect of this re-scaling factor and show that taking a value of $\xi = 10$ makes only minimal difference in the derived limits, thus justifying the adopted re-scaling factor used for the fiducial analysis.

3.B Pulsar data

In order to constrain the existence of axions, we compare the predicted radio flux with observations from 27 nearby pulsars. We have limited ourselves to this small, carefully selected subsample in order to make the analysis computationally feasible. Our selection is largely based on two fundamental criteria: we want pulsars with large magnetic fields and pulsars with a well-measured orientation (meaning that both the misalignment angle and the pulsar’s orientation relative to Earth are known). The former is particularly important as the production of axions scales $\propto B_s^4$ and the resonant conversion process scales $\propto B(\vec{x}_c)^2$ (with \vec{x}_c being the point of resonant conversion). We have chosen to apply the geometrical constraint as the axion-induced radio flux varies rather strongly with viewing angle (see *e.g.* Figure 3.2), and knowledge of the geometry evades the need for marginalization. It is worth mentioning that a broader study over the entire pulsar population would naturally marginalize over the geometry, and could yield stronger constraints than those obtained here; currently, such a study is computationally prohibitive, however we hope that future improvements will make this manageable.

The radio observations were obtained using the Low-Frequency Array [295], the Lovell telescope at Jodrell Bank [296], and the Shanghai Tian Ma Radio Telescope [297], and span a wide range of frequencies from 25 MHz to 8.6 GHz¹⁰. The properties of each of the 27 pulsars and the observed flux densities are listed for reference in Tables 3.1 and 3.2. In the next sections, we outline the analysis used to derive upper limits on the axion-photon coupling, and highlight the impact of various systematic uncertainties in the modeling of the axion signal.

¹⁰The low frequency spectrum of radio pulsars can, in some cases, be suppressed by free-free absorption [298]. For each of the pulsars of interest, we use the observed dispersion measure to infer the average line of sight electron number density, and subsequently use this quantity to estimate the relevant optical depth. For all pulsars of interest we find that the optical depth $\tau \ll 1$, suggesting free-free absorption can be neglected.

Name	Period (s)	B_s (G)	θ_m (deg)	β (deg)	Distance (kpc)
B1822-09	0.77	6.42E12	95	-7	$0.3^{+0.7}_{-0.2}$
B2045-16	1.96	4.69E12	36	1.1	$0.95^{+0.02}_{-0.03}$
B0450-18	0.55	1.80E12	24	4	$0.4^{+0.2}_{-0.1}$
B0329+54	0.71	1.22E12	30	2.1	$1.0^{+0.1}_{-0.1}$
B0450+55	0.34	9.10E11	32	3.3	$1.18^{+0.07}_{-0.05}$
B0656+14	0.38	4.65e12	30	8.2	$0.29^{+0.03}_{-0.03}$
B0919+06	0.43	2.46e12	48	4.8	$1.1^{+0.2}_{-0.1}$
B1508+55	0.74	1.95e12	45	-2.7	$2.1^{+0.1}_{-0.1}$
B1738-08	2.04	2.18E12	26	1.7	$0.2^{+0.2 \dagger}_{-0.2}$
B1818-04	0.60	1.97E12	65	3.5	$0.6^{+0.03}_{-0.03}$
B1917+00	1.27	3.16E12	81	1.3	$3.6^{+0.5 \dagger}_{-0.5}$
B1919+14	0.62	1.88E12	26	-6.4	$3.4^{+0.55 \dagger}_{-0.55}$
B1919+21	1.34	1.36E12	45	-3.7	$0.3^{+0.8}_{-0.2}$
B1920+21	1.08	3.00E12	44	1.1	$4.0^{+2.0}_{-2.0}$
B2224+65	0.68	2.60E12	16	3.4	$0.9^{+0.1}_{-0.1}$
B2319+60	2.26	4.03E12	18	2.2	$2.7^{+1.2}_{-0.9}$
B1910+20	2.23	4.82E12	29	1.5	$3.6^{+0.3 \dagger}_{-0.3}$
B0402+61	0.59	1.84E12	83	2.2	$1.95^{+0.17 \dagger}_{-0.17}$
B1039-19	1.39	1.16E12	31	1.7	$2.0^{+0.5 \dagger}_{-0.5}$
B1237+25	1.38	1.17E12	53	0	$0.84^{+0.06}_{-0.05}$
B1737+13	0.80	1.09E12	41	1.9	$4.2^{+4.2 \dagger}_{-4.2}$
B1737-30	0.61	1.70E13	58	10.9	$0.4^{+1.7}_{-0.3}$
B0833-45	0.089	3.38E12	90	12.0	$0.28^{+0.02}_{-0.02}$
B0923-58	0.74	1.93E12	19	3.0	$0.107^{+0.1 \dagger}_{-0.1}$
B1742-30	0.367	2.0E12	24	6.4	$0.2^{+1.1}_{-0.2}$
B1749-28	0.56	2.16E12	42	2.9	$0.2^{+1.1}_{-0.1}$
B2334+61	0.495	9.91E12	33	3.5	$0.7^{+0.1}_{-0.1}$

Table 3.1: List of pulsars used in this chapter. For each pulsar the columns denote: name, rotational period (in seconds), inferred dipolar surface magnetic field B_s (in Gauss), the angle θ_m between the magnetic and rotational axis (in degrees), the angle β between the magnetic axis and the observer at closest approach (in degrees), and the distance (in kpc). The inferred geometric angles are obtained from [299–301]. Distance measurements and errors are taken from [302–305]. Following the ATNF procedure, we prioritize distance measurements not inferred using the dispersion measure and galactic electron models. In some cases (marked with a †), the galactic coordinates and dispersion measure are used to infer the distance utilizing the online tools from [305]. Should the various galactic models agree, we take the mean and standard deviation of these inferred distances; if they instead differ by more than a factor of two, we defer to the more recent model of [303] and place a 100% error on the distance.

Name	S_{25-80} (mJy)	S_{50-80} (mJy)	S_{408} (mJy)	S_{1400} (mJy)	S_{8600} (mJy)
B1822-09	-	2502 ± 1251	35.5 ± 2.5	10.2 ± 20	-
B2045-16	-	-	115 ± 26	22 ± 44	0.3 ± 0.04
B0450-18	-	-	82.4 ± 10	16.8 ± 34	-
B0329+54	-	1841 ± 921	1512 ± 175	203 ± 57	-
B0450+55	124 ± 62	-	59.0 ± 9	13.0 ± 3	-
B0656+14	-	-	6.5 ± 0.6	2.7 ± 2	0.96 ± 0.11
B0919+06	-	550 ± 275	52.2 ± 6.4	10 ± 3	-
B1508+55	-	943 ± 471	114 ± 5	8 ± 1	-
B1738-08	-	-	29.4 ± 7.7	1.4 ± 4.0	-
B1818-04	-	-	156 ± 6	10.07 ± 2.0	-
B1917+00	-	-	15.6 ± 1.3	0.8 ± 2.0	-
B1919+14	-	-	< 5.2	0.68 ± 8	-
B1919+21	1586 ± 793	-	56.5 ± 8.2	18.8 ± 38	-
B1920+21	-	-	29.9 ± 1.2	1.4 ± 2.0	-
B2224+65	-	293 ± 146	21.9 ± 2.4	2 ± 4	-
B2319+60	-	-	36.1 ± 4.9	12 ± 1	-
B1910+20	-	-	5.7 ± 1.1	0.22 ± 3	-
B0402+61	-	-	14.6 ± 1.3	2.8 ± 2	-
B1039-19	-	-	27.7 ± 5.5	0.62 ± 7	-
B1237+25	102 ± 51	-	110 ± 33	23.2 ± 47.0	0.31 ± 0.07
B1737+13	-	131 ± 66	23.9 ± 1.8	8.9 ± 5	-
B1737-30	-	-	24.6 ± 0.1	8.9 ± 5	1.21 ± 0.07
B0833-45	-	-	5000 ± 0.1	1050 ± 60	-
B0923-58	-	-	22 ± 0.1	21 ± 6	-
B1742-30	-	-	66 ± 6	21 ± 1	1.55 ± 0.05
B1749-28	-	-	1100 ± 100	47.8 ± 96	0.3 ± 0.05
B2334+61	-	-	10.0 ± 2	1.4 ± 3	-

Table 3.2: List of pulsars used in this chapter. For each pulsar the columns denote: name, the flux density between 25 – 80 MHz [295], the flux density between 50 – 80 MHz [295], the flux density at 408 ± 4 MHz [296], the flux density at 1.4 ± 0.032 GHz [296], and the flux density at 8.6 ± 0.8 GHz [297] (all in mJy).

3.C Profile likelihood analysis

We utilize the profile likelihood ratio to derive a one-sided upper limit on $g_{a\gamma\gamma}$. In general, this is accomplished by defining a test statistic [306–308]

$$\text{TS}(\mu) = \begin{cases} -2 \ln \frac{\mathcal{L}(\mu, \hat{\vec{\theta}})}{\mathcal{L}(\hat{\mu}, \hat{\vec{\theta}})} & \text{if } \hat{\mu} \leq \mu, \\ 0 & \text{if } \hat{\mu} > \mu. \end{cases} \quad (3.33)$$

Here μ is the parameter of interest (in our case, $g_{a\gamma\gamma}$), $\vec{\theta}$ denotes the nuisance parameters (which here include the pulsar distances and the intrinsic radio flux from each pulsar at each frequency), $(\hat{\mu}, \hat{\vec{\theta}})$ are the maximum likelihood estimators, and $\hat{\vec{\theta}}$ is the conditional maximum likelihood estimator for a given value of μ .

We adopt Gaussian likelihoods for the flux density and pulsar distances. The full likelihood is given by

$$\mathcal{L}(g_{a\gamma\gamma}, \vec{\theta}) = \prod_i^{N_{\text{pul}}} \prod_j^{N_{\text{bin}}} e^{-(S_{\text{obs},ij} - S_{\text{pred},ij})^2 / 2\sigma_{S,ij}^2} \times e^{-(d_i - \mu_{d,i})^2 / 2\sigma_{d,i}^2}, \quad (3.34)$$

where the product is taken over each pulsar and each frequency bin for which observations are available (see Tables 3.1 and 3.2). Some important comments regarding this equation:

- The predicted flux density has two components, one due to the predicted axion signal and another that is intrinsic to the respective pulsar, *i.e.* $S_{\text{pred},ij} = S_{\text{axion},ij} + S_{\text{pulsar},ij}$. Owing to the fact that we do not currently have a firm understanding of the intrinsic radio flux produced by pulsars, we treat the parameters $S_{\text{pulsar},ij}$ as nuisance parameters. This implies that our test statistic is zero unless the predicted axion flux in at least one observing bin of one pulsar exceeds the observed flux. This procedure ensures that our constraints remain conservative.
- All observed flux densities and mean distances, as well as the variances in both of these quantities, are reported in Tables 3.1 and 3.2. Note that the variance in the distance is generally asymmetric, meaning that the second exponential in Equation 3.34 is actually an asymmetric Gaussian.
- The pulsar distances show up explicitly in the likelihood, but also enter through the predicted flux densities.

With the likelihood in hand, we compute our test statistic at set couplings and derive an upper limit on $g_{a\gamma\gamma}$ by identifying the value at which $\text{TS} = 2.71$. This corresponds to the 95% upper limit [308].




Figure 3.8: Fiducial upper limits derived for the simulation (solid) and semi-analytic model (dotted), shown together with corresponding bands depicting systematic model uncertainty.




Figure 3.9: *Left:* Upper limits derived from one of the strongest pulsars (in the semi-analytic model) using exclusively observations of low (~ 408 MHz), mid (~ 1.4 GHz), or high (~ 8.6 GHz) frequencies. *Right:* Comparison of the limits derived from three of the strongest pulsars using both the simulation (solid) and the semi-analytic model (dotted). The bounds for each pulsar differ by no more than a factor of a few between models.

The derived limits for the simulation and semi-analytic model are shown in Figure 3.8, where the lines represent the fiducial constraints and the bands reflect the uncertainty due to particular modeling choices – the bands are also shown in Figure 3.3 and discussed thoroughly in the next section. We find that the final limits obtained through our pipeline are largely driven by observations of a few pulsars at higher frequencies. The reason for the latter is that the pulsar flux tends to drop much faster

than the predicted axion flux with increasing frequency. In order to illustrate the origin of the constraining power, we plot in Figure 3.9 the constraints derived from one of the strongest pulsars (in the semi-analytic model) using exclusively observations of low (~ 408 MHz), mid (~ 1.4 GHz), or high (~ 8.6 GHz) frequencies (left panel). We also plot in the right panel of Figure 3.9 the individual constraints obtained from both models using the three strongest pulsars (utilizing data at all frequencies). We find that the relative differences between the two models is typically no larger than a factor of ~ 2 in coupling (although larger differences appear for a small amount of outliers, like pulsar B0656+14 at intermediate masses), providing confidence that reasonable changes to our axion production model have minimal impact on the final constraints.

3.D Uncertainties

This section reports the various uncertainties that enter our calculations, including uncertainties associated with the inferred pulsar properties, uncertainties in treating the PIC simulation, and uncertainties associated with the free parameters entering the semi-analytic model. We also provide quantitative estimates of the effects due to these uncertainties on the final constraints.

3.D.1 Observational uncertainties

Let us begin by discussing uncertainties associated with the inferred properties of each pulsar. In general, the properties that are measured to high precision for every pulsar are the rotational period P_{NS} and the spin-down rate \dot{P}_{NS} . The value of the surface dipolar magnetic field strength B_s can be inferred from these quantities. This can be seen by comparing the net Poynting flux through the light cylinder of a misaligned rotator, given by $S = \mu^2 \Omega_{\text{NS}}^4 (1 + \sin^2 \theta_m)$ (where $\mu = B_s R_{\text{NS}}^3$, $\Omega_{\text{NS}} = 2\pi/P_{\text{NS}}$, and θ_m is the misalignment angle [309]), to the spin-down luminosity, yielding

$$B_s = \frac{1}{2\pi} \left(\frac{I_{\text{NS}}}{R_{\text{NS}}^6 (1 + \sin^2 \theta_m)} \right)^{1/2} (P_{\text{NS}} \dot{P}_{\text{NS}})^{1/2}. \quad (3.35)$$

There are two points worth noting here. First, the ATNF adopts a fiducial moment of inertia $I_{\text{NS}} = 10^{45}$ g cm², which is actually 30 – 50% lower than current estimates [310]. Second, the ATNF inferred field strengths assume $\theta_m = 0$. One can see that these effects somewhat offset, however likely lead to a slight underestimation of B_s at the level of $\sim \mathcal{O}(10\%)$.

Another important consideration is that Equation 3.35 only reflects the magnitude of the dipolar magnetic field – near the surface of the neutron star additional higher order multipoles could, and in some cases are expected to, contribute. This likely implies that the values of B_s listed in Table 3.1 underestimate the true surface magnetic

field strength (and thus we are likely underestimating our signal, potentially by a significant amount). Incorporating higher multipoles, however, must be done self-consistently, and thus in what follows we adopt the assumption that the magnetic fields are purely dipolar. We believe that this is likely a conservative assumption, which we plan to address more rigorously in future work.

Since the derived limit is expected to scale like $g_{a\gamma\gamma}^{\text{lim}} \propto B_s^\alpha$ with $1 \lesssim \alpha \lesssim 1.5$ (depending on the efficiency of the resonance), we adopt in what follows a characteristic uncertainty on $g_{a\gamma\gamma}^{\text{lim}}$ of $\pm 20\%$. We expect this to be conservative, as it underestimates the impact of the corrected moment of inertia and neglects the contribution of higher order multipoles, both of which are expected to enhance the signal.

Notice that in estimating the uncertainty in the dipolar field strength we have neglected the contribution coming from the unknown equation of state, which allows for $\sim 10\%$ fluctuations in R_{NS} [311]. This is because the neutron star radius enters at multiple points in the analysis, and thus we treat this as an independent uncertainty. In particular, we expect the neutron star radius to enter with the following scalings: $r_{\text{pc}} \propto R_{\text{NS}}^{3/2}$, $E_{\parallel} \propto r_{\text{pc}} \propto R_{\text{NS}}^{3/2}$, and $B_s \propto R_{\text{NS}}^{-3}$. Roughly speaking, this translates into a scaling on the limits as $g_{a\gamma\gamma}^{\text{lim}} \propto (B_s^{4\alpha} E_{\parallel}^2 r_{\text{pc}}^4)^{-1/4} \propto R_{\text{NS}}^{9/4}$, where as before $1 \lesssim \alpha \lesssim 1.5$. Thus, a 10% variation in the neutron star radius amounts to a 20% shift in the limits.

There exist a variety of techniques for inferring the distances to observed pulsars. By far the most reliable of these are parallax measurements, but unfortunately such measurements are only available for a small number of known pulsars. One can also analyze the HI spectrum on- and off-pulse, and use this information to infer the characteristic rotational velocities of the HI regions behind and in front of the pulsar. In many cases, however, this information is not available. Without additional information, one must rely on inferring the distance using the observed dispersion measure ($\text{DM} = \int dl n_e$), which requires a model for the galactic electron density. In Table 3.1 we compile inferred distances and uncertainties using various combinations of the techniques listed above. Priority is given to geometric measurements, and we only rely on dispersion measure distances when no other information is available (as done in the ATNF) [302–305]. In some cases, no published distance could be found; here, we use a variety of galactic electron density models to find the distances using the tools available in [305]. In a few cases, uncertainties on the distance were not readily available. In these cases, we compared the inferred distance from two different galactic electron density models – if the models showed reasonable agreement (differing by no more than a factor of 2), we inferred the standard deviation from these measurements. In all other cases we defaulted to the more recent model [303] and placed a 100% uncertainty on the distance. It is worth emphasizing that none of these pulsars (denoted in Table 3.1 with a †) contribute meaningfully to the derived constraints.

Finally, let us mention that the uncertainties on the pulsar geometry (the misalignment angle and orientation) as interpreted within the rotating vector model (RVM) [312] for the pulsars used in this analysis are expected to be small (see *e.g.* [300]). Note that this is not a generic feature of all pulsars, however, we have chosen our pulsars largely to minimize this uncertainty. There is of course a larger systematic uncertainty that pertains to the validity of the model itself, but it is unclear how this can be quantified. Given that the RVM is the state-of-the-art, we take these values to be fixed in our analysis – future analyses which are applied over large populations of pulsars may be able to evade the need for specific assumptions on pulsar geometry. We leave such an endeavor for future work.

3.D.2 PIC simulation

There are two clear identifiable uncertainties that enter our analysis of the PIC simulation. First, the physical dimensions of the simulation are expressed in terms of the polar-cap radius. In general, the fractional surface area which leads to pair cascades depends on a variety of factors including the magnetic field geometry, the misalignment angle, the properties of the return current, the compactness of the neutron star, and various non-linearities in the electrodynamics. While it is difficult to quantify the effects of each of these separately, it is expected that together they can lead to an $\mathcal{O}(1)$ modification of the polar-cap radius [188, 313, 314]¹¹. Axion production at low masses scales with the square of the volume (with a weakening of the dependence on the volume at higher masses), meaning we expect our limit on the axion-photon coupling to scale proportional to this uncertainty. In what follows we therefore adopt a characteristic uncertainty on r_{pc} at the level of $\pm 30\%$ ¹², which roughly corresponds to the one-sigma range inferred from a flat distribution varying between $\pm 50\%$ of the fiducial value.

The second source of uncertainty comes from the re-scaling factor ξ , which is intended to undo the compression of scales that is performed in the simulation. Unfortunately, without more simulations at various resolutions it is impossible to truly understand the impact of this re-scaling factor. Given, however, that the semi-analytic model also contains a slight scale compression (and subsequent re-scaling), we adopt the uncertainty inferred from the semi-analytic model in the simulation as well (see following subsection). We do note that this procedure is not exactly valid given that

¹¹Note that pulsars with extremely large non-dipole multipoles can have smaller polar-cap radii. This however comes at the expense of an enormous magnetic field, which is expected to entirely offset the resultant decrease in axion production [315].

¹²Note that the PIC simulation intrinsically assumes an aligned rotator (oblique rotators would require going to 3+1 dimensions). The physics of pair production and screening is expected to be the same for oblique rotators, with the dominant effect being a shift in the geometry and the potential drop – this, however, is an implicit assumption which must be verified in the future.

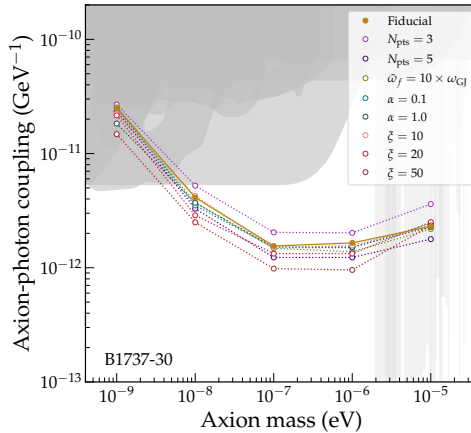


Figure 3.10: Upper limits derived using observations of B1737-30 and the semi-analytic model for the axion production rate. We assess the impact of varying several unconstrained parameters (namely N_{pts} , the final value of $\tilde{\omega}$, α , and ξ) on the derived limits, and compare with the fiducial model (solid, gold).

the scale compression is not done in an equivalent manner – it is largely for this reason that the semi-analytic model was constructed in the first place.

Finally, it is worth mentioning that there also exists an unquantifiable uncertainty that enters the analysis of the PIC simulation, which arises from the fact that the physical parameters used in the simulation do not directly correspond to those of true pulsars (recall that this is because the parameter re-scalings performed in the simulation modify the fundamental scales of the problem). In the future, one would ideally like to perform a number of simulations varying the re-scaling parameters in order to understand the sensitivity of axion production to this procedure, however at the moment this is computationally unfeasible. Thus, we instead use the semi-analytic model as a cross-check, allowing us to ensure the robustness of the results.

3.D.3 Semi-analytic model

As with the simulation, the semi-analytic model is subject to uncertainties in the characteristic polar-cap size. Additionally, the semi-analytic model has a number of free parameters which have been fixed to fiducial values in the derivation of the upper limits. Here, we outline the impact of reasonable variations in each of these. The primary free parameters include: N_{pts} (the number of seed points used in the plane wave damping), $\tilde{\omega}_f$ (the final plasma frequency at the end of the particle cascade), α (the power-law index controlling how quickly the plasma frequency evolves after the initial collapse), ξ (the re-scaling factor), ϵ (the parameter controlling the acceleration time of particles in the gap), and κ (the parameter describing the initial plasma density

Parameter	Characteristic uncertainty on $g_{a\gamma\gamma}^{\text{lim}}$	Simulation	Semi-analytic
B_s	$\pm 20\%$	✓	✓
R_{NS}	$\pm 20\%$	✓	✓
r_{pc}	$\pm 30\%$	✓	✓
N_{pts}	$\pm 25\%$	✗	✓
$\hat{\omega}_f$	$\pm 5\%$	✗	✓
α	$\lesssim \pm 5\%$	✗	✓
ξ	$\lesssim \pm 5\%$	✓	✓
ϵ	$\sim 0\%$	✗	✓
κ	$\sim 0\%$	✓	✓
Total		$\pm 42\%$	$\pm 49\%$

Table 3.3: Estimated breakdown of the characteristic uncertainty on the derived limits arising from systematic uncertainties in the semi-analytic model and simulation (we denote in each column whether the listed uncertainty is included and relevant for each model). The total uncertainty is obtained by combining the individual contributions in quadrature, and is reflected in the final row of each model.

present in the gap prior to the pair cascade – note that this parameter also enters the simulation).

The latter two parameters (namely, ϵ and κ) have no perceptible effect on the final results. In the semi-analytic model, the acceleration time (see Equation 3.17) is the smallest length scale in the problem. This means that for the fields considered in this chapter, the acceleration time is often smaller than the time resolution of the model. As a result, we generally accelerate each particle to its final Lorentz factor within a single time step. The value of ϵ , which describes at what fraction of the radiation-reaction-limited Lorentz factor each particle begins to emit CR photons, shows up logarithmically in the acceleration time and thus has at most an $\mathcal{O}(1)$ effect on t_{acc} , which is already small compared to the time resolution. Therefore, even orders of magnitude changes in ϵ have negligible effect on the final result. The parameter κ is introduced to avoid having to simulate $\sim n_{\text{GJR}}^3 r_{\text{pc}}^3$ particles, which would be computationally infeasible. Since the number density grows exponentially during a cascade, κ can be understood as a logarithmic re-scaling of the time in the pair cascade process. However, we note that the pair cascade process is introduced to simply set the initial conditions in the gap prior to dynamical screening. The initial conditions and the dynamics of the screening phase are independent of κ , and thus the effect of κ on the final result is also negligible.

In order to assess the sensitivity of the presented limits to the choice of the other parameters we run an independent analysis on one of the most constraining pulsars, B1737-30. In particular, we perform an extra run with $N_{\text{pts}} = 3$, one with $N_{\text{pts}} = 5$

(rather than $N_{\text{pts}} = 4$), a run in which $\hat{\omega}_f = 10$ (rather than 100), two runs in which α is varied from 0.1 to 1.0, and various runs using a larger stretching parameter ξ . The limits obtained in each of these runs are displayed alongside our fiducial limit from this pulsar in Figure 3.10. The characteristic uncertainty of each parameter variation on the limit is quantified and included in Table 3.3. Note that in the special case of ξ , we infer the uncertainty by looking at the limit in which $\xi \rightarrow 1$ rather than by comparing variations across runs.

Most of the uncertainties listed in Table 3.3 are extremely subdominant to that coming from the uncertainty associated to the polar-cap size, with the variation in N_{pts} otherwise producing the largest effect. The total uncertainty for each model (estimated at the level of 42% and 49% for the simulation and semi-analytic models, respectively) is obtained by summing the errors in quadrature, and is reflected in a band about the fiducial limits in Figure 3.3 and Figure 3.8.