

Constraints on light QCD and CP -violating axions from the death line of rotation-powered pulsars

Samuel J. Witte^{1,2,3,*}, Andrea Caputo^{4,5,6,†}, Stefan Stelzl^{7,8,‡}, Alexander Chernoglazov⁹,
Alexander A. Philippov¹⁰, and Surjeet Rajendran¹¹

¹*Rudolf Peierls Centre for Theoretical Physics, University of Oxford, Oxford, United Kingdom*

²*Deutsches Elektronen-Synchrotron DESY, Notkestraße 85, 22607 Hamburg, Germany*

³*II. Institute of Theoretical Physics, Universität Hamburg, 22761, Hamburg, Germany*

⁴*Department of Theoretical Physics, CERN, Esplanade des Particules 1,
P.O. Box 1211, Geneva 23, Switzerland*

⁵*Dipartimento di Fisica, “Sapienza” Università di Roma and Sezione INFN Roma 1,
Piazzale Aldo Moro 5, 00185 Roma, Italy*

⁶*Department of Particle Physics and Astrophysics, Weizmann Institute of Science, Rehovot 7610001, Israel*

⁷*Institute of Physics, Theoretical Particle Physics Laboratory,
Ecole Polytechnique Fédérale de Lausanne, Switzerland*

⁸*Institut de Física d’Altes Energies (IFAE) and Barcelona Institute of Science and Technology (BIST),
Campus UAB, 08193 Bellaterra (Barcelona), Spain*

⁹*Institute for Advanced Study, Princeton, New Jersey 08540, USA*

¹⁰*Department of Physics, University of Maryland, College Park, Maryland 20742, USA*

¹¹*Department of Physics and Astronomy, Johns Hopkins University, Baltimore, Maryland 21218, USA*



(Received 9 February 2026; accepted 4 August 2026; published 28 September 2026)

Dense nuclear matter can modify the effective potential of axions, displacing them from their vacuum minimum and sourcing large external field gradients (“axion hair”). In the case of neutron stars, axion hair directly modifies the electrodynamic processes operating on the open field-line region, strongly enhancing or suppressing the acceleration experienced by ambient charges. As a result, the point in the neutron star lifetime at which pair cascades cease—known as pulsar “death”—can be dramatically altered, allowing for much older pulsars to emit observable radio emission. We study the pair discharge process in the presence of axion hair using semianalytic techniques and particle-in-cell simulations and use these results alongside pulsar demographics to derive new constraints on light QCD axions with non-negligible axion-photon coupling and CP -violating axion-nucleon interactions. We also illustrate how nearly orthogonal rotators, where emission is observed from both poles (such as in the case of PSR J1906 + 0746), provide a complementary probe of axion hair.

DOI: [10.1103/jdmb-rws8](https://doi.org/10.1103/jdmb-rws8)

I. INTRODUCTION

The quantum chromodynamics (QCD) axion is a pseudo-Nambu-Goldstone boson arising from the spontaneous breaking of the Peccei-Quinn symmetry [1–3] and provides one of the most compelling solutions to the strong CP problem (i.e., the question of why the neutron electric dipole moment is so small [4]). In typical UV-complete models, the QCD axion is expected to live

in a well-defined region of parameter space; this is because the effective low-energy potential generated by QCD relates the axion mass to the decay constant, $m_a \sim 5.7 \mu\text{eV}(10^{12} \text{ GeV}/f_a)$, where f_a also sets the characteristic size of the axion couplings. This relation, however, is not fully generic: there exist mechanisms in which the axion remains parametrically lighter than in conventional scenarios [5–8]. Such models are especially interesting from a phenomenological perspective, since localized dense clumps of matter can distort the effective axion potential, displacing the axion from the vacuum minimum and sourcing large field gradients within—and near the boundaries of—these dense objects. Matter-sourcing effects for light axions have been studied in a variety of contexts ranging from Earth and the Sun [7] to white dwarfs [9–11] and neutron stars (NSs) [7,12–15]. NSs represent one of the most interesting scenarios (albeit also one of the most complicated), as the extreme densities in

*Contact author: Samuel.Witte@physics.ox.ac.uk

†Contact author: Andrea.Caputo@cern.ch

‡Contact author: SStelzl@ifae.es

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article’s title, journal citation, and DOI. Funded by SCOAP³.




FIG. 1. Rotational period P and spin-down rate $\dot{P}$ of the rotation-powered pulsar population (blue). Two standard estimates of the pulsar death line, which differentiate the electromagnetically active from the inactive population, are shown in yellow. The red line shows the shift in the death line induced by a light QCD axion with $\epsilon = 10^{-2}$ and $f_a = 10^{16}$ GeV. The construction of the death-line curves and the discharge modeling underlying them are presented in the companion paper [20].

their interiors allow axion gradients to be sourced across a broader region of parameter space.

In this work, we propose a novel pulsar signature arising from static axion field gradients, or “axion hair,”¹ near NSs. We demonstrate that axion hair can significantly modify the electrostatic potential drop along the open magnetic field lines of rotation-powered pulsars, allowing pair production and coherent radio emission to occur in much older, more slowly rotating NSs than would otherwise be possible. In effect, axion hair can extend pulsar lifetimes by decoupling the acceleration of primary $e^\pm$ from the rotational period of the pulsar itself, generating a shift in the location of the so-called pulsar “death line” [17–19]. We argue that the lack of observation of such an old isolated pulsar population cannot be explained by observational selection effects alone, leading to an apparent incompatibility with the presence of sufficiently long-range axion hair sourced by light axions coupled to nuclear matter. This picture is illustrated in Fig. 1, where axion hair sourced by a representative light QCD axion induces a large shift in the pulsar death line.

A second, qualitatively different observable follows from the intrinsic “polarity asymmetry” introduced by axion hair: for sufficiently long-range gradients, the two magnetic hemispheres can be driven into different discharge regimes,

¹This terminology was also used in [16] to describe local axion gradients sourced near NSs.

which can be incompatible with a rare class of objects in which coherent radio emission is observed twice per period (so-called interpulses). We validate our arguments using particle-in-cell (PIC) simulations of the near-surface discharge; the companion article [20] provides a complete description of pulsar electrodynamics in the presence of axion hair, including our numerical setup and an extended treatment of the discharge physics, global consistency requirements, and systematics.

We also demonstrate how both effects arise in models where axions have, in addition to the coupling to photons, a linear coupling to nucleons (leading to a monopole-dipole interaction) [21]; in this case, interpulses observed in PSR J1906 + 0746 can be used to derive constraints comparable to those obtained using the death line. For this class of axions, both CP and the axion shift symmetry are explicitly broken² (e.g., as in the QCD axion with a nonzero θ angle [22,23]). It is worth noting that CP -violating nucleon interactions are typically severely constrained by fifth-force experiments [12,24], highlighting the strength and complementarity of the pulsar-based effects studied here.

II. SOURCING AXIONS FROM NUCLEAR MATTER

We begin by reviewing how light QCD axions are sourced in nuclear matter, before commenting on the monopole-dipole interaction [7,9,12,25]. Additional details (including sourcing regimes, phase structure, and stellar-structure/equation-of-state systematics) are discussed in the companion paper [20].

Consider a pseudoscalar field coupled to gluons,

$$\mathcal{L} \supset \frac{1}{2}(\partial a)^2 + \frac{g_s^2}{32\pi^2} \frac{a}{f_a} G\tilde{G}, \quad (1)$$

where a is the axion, f_a the decay constant, g_s the strong coupling, and $G_{\mu\nu}$ the gluon field strength. Below the QCD confinement scale, nonperturbative dynamics generate an effective potential [26,27],

$$V(a) = -\epsilon m_\pi^2 f_\pi^2 \left(\sqrt{1 - z_{\text{ud}} \sin^2 \left(\frac{a}{2f_a} \right)} - 1 \right), \quad (2)$$

where $z_{\text{ud}} = 4m_u m_d / (m_u + m_d)^2$. The conventional QCD axion corresponds to $\epsilon = 1$, while symmetry-based models producing lighter QCD axions with $\epsilon < 1$ are discussed in [5–8,28,29]. Although the mass around the origin remains

²This could also occur for a CP -conserving nucleon coupling should the star be polarized. Here, we have focused the CP -odd coupling to axions, but one could also consider a CP -even coupling, which would induce a shift in the fine structure constant; typically, however, these constraints are much more stringent than their CP -odd counterparts.

unchanged, some of these models differ from Eq. (2) in their vacuum structure; we discuss model-building variants in the companion paper [20].

The axion mass is given by

$$m_a = \sqrt{\epsilon} \frac{\sqrt{z_{ud}}}{2} \frac{m_\pi f_\pi}{f_a} \simeq 5.8 \times 10^{-7} \text{ eV} \left(\frac{\epsilon}{10^{-2}} \right)^{1/2} \left(\frac{10^{12} \text{ GeV}}{f_a} \right). \quad (3)$$

At low energies, a shift-symmetry-breaking interaction between axions and nucleons is generated,

$$\mathcal{L} \supset \sigma_{\pi N} \bar{N} N \left(\sqrt{1 - z_{ud} \sin^2 \left(\frac{a}{2f_a} \right)} - 1 \right), \quad (4)$$

where $\sigma_{\pi N} \simeq 50 \text{ MeV}$ is the pion-nucleon σ term. In a background of nonrelativistic nucleons with number density n one finds $\langle \bar{N} N \rangle \simeq n \equiv \langle \bar{N} \gamma^0 N \rangle$. The above interaction leads to a new term in the effective axion potential, and if

$$\sigma_{\pi N} n_b \gtrsim \epsilon m_\pi^2 f_\pi^2 \Rightarrow n_b \gtrsim 0.16 \text{ fm}^{-3} \left(\frac{\epsilon}{0.4} \right), \quad (5)$$

the minimum of the effective potential shifts from $a = 0$ to $a = \pi f_a$. Equivalently, the interaction in Eq. (4) can be interpreted as an axion-dependent nucleon mass.

Equation (5) is a necessary, but not sufficient, condition for sourcing: the system must also be large enough to accommodate the axion gradient. Sourcing occurs if the system size $R \gg m_a(n)^{-1}$, where $m_a(n)$ is the effective in-density axion mass (note that we reserve the notation of m_a to represent the axion vacuum mass). In NSs this condition is satisfied for $f_a \lesssim 10^{17} \text{ GeV}$ [12].

Depending on the value of ϵ , there are two different regimes. For $\epsilon \lesssim 0.07$, the axion is sourced all the way until the boundary of the star, while for $\epsilon \gtrsim 0.07$, a phase transition between the axion-sourced phase and the normal nuclear phase happens inside the star. For this paper, we focus primarily on the first regime (see the companion paper [20] for a detailed discussion), where the axion profile [12] is well approximated by

$$a(r) = \begin{cases} \pm \pi f_a, & r < R, \\ \pm \pi f_a \frac{R}{r} \exp \left[-\frac{\sqrt{\epsilon z_{ud}} m_\pi f_\pi}{2f_a} (r - R) \right], & r > R, \end{cases} \quad (6)$$

with characteristic decay length

$$\lambda \simeq \frac{f_a}{\sqrt{\epsilon} m_\pi f_\pi} \simeq 16 \text{ cm} \left(\frac{f_a}{10^{12} \text{ GeV}} \right) \left(\frac{10^{-2}}{\epsilon} \right)^{1/2}. \quad (7)$$

The example above applies explicitly to the light QCD axion, but similar results may arise generically for other light scalars. Consider, for example, a scalar linearly coupled to

nucleons, $\mathcal{L} \supset \frac{1}{2}(\partial a)^2 - \frac{1}{2}m_a^2 a^2 + g_N a \bar{N} N$. In this case, the minimum of the potential shifts smoothly with ambient density,

$$a_{\min} = g_N \frac{\langle \bar{N} N \rangle}{m_a^2}. \quad (8)$$

This continuous shift, however, only happens on length scales longer than m_a^{-1} . On shorter length scales, the axion gradient stops the field from tracking its minimum. Thus there is an axion gradient leaking out of the star. Unlike for the QCD axion, the viable parameter space typically requires $a \ll m_N/g_N$ inside NSs [12], implying a negligible backreaction on the stellar structure and highlighting the importance of the electrodynamic effects studied here.

III. PULSAR ELECTRODYNAMICS AND THE IMPLICATIONS OF AXION HAIR

Having demonstrated how axion hair arises, we turn our attention to the impact of these field gradients on the electromagnetic evolution of NSs.

As a result of the coupling of the axion to quarks and mesons, QCD axions acquire a coupling to electromagnetism via

$$\mathcal{L} \supset -\frac{g_{a\gamma\gamma}}{4} a F_{\mu\nu} \tilde{F}^{\mu\nu}, \quad (9)$$

where $g_{a\gamma\gamma} = \alpha_{\text{EM}} C / (2\pi f_a)$ and C is an $\mathcal{O}(1)$ model-dependent coefficient. This interaction modifies Gauss's and Ampère's laws,

$$\nabla \cdot \vec{E} = \rho - g_{a\gamma\gamma} \vec{B} \cdot \nabla a \equiv \rho - \rho_a, \quad (10)$$

$$\nabla \times \vec{B} - \partial_t \vec{E} = \vec{j} + g_{a\gamma\gamma} \dot{a} \vec{B} + g_{a\gamma\gamma} \nabla a \times \vec{E}. \quad (11)$$

Axion hair sourced by stellar objects is static, $\partial_t a \simeq 0$, and the rotationally induced electric field near the surface of the star is always much less than the magnetic field, i.e., $|\vec{E}|/|\vec{B}| \ll 1$, implying the leading correction comes from $\vec{B} \cdot \nabla a$. Comparing the effective axion charge density near the stellar surface, $\rho_a \equiv g_{a\gamma\gamma} \vec{B} \cdot \nabla a$, with the corotating (Goldreich-Julian) charge density $\rho_{\text{GJ}} \simeq -2\vec{\Omega} \cdot \vec{B}$, one sees that, for pulsars near the conventional death line, $|\rho_a/\rho_{\text{GJ}}| \sim (1 + m_a R) \alpha_{\text{EM}} / (R\Omega) \gg 1$. Thus, on distances $r \lesssim m_a^{-1}$, axion hair rather than rotation controls the local charge density requirement and can dominate the near-field electrodynamics.

A. Pulsar death

We briefly review the physics behind the pulsar “death line,” which differentiates electromagnetically active from inactive pulsars [17–19,30–34].

For typical rotation-powered pulsars, the dominant coherent radio emission is produced near the star on the

open field-line bundle. The precise emission mechanism has been investigated for decades, with recent progress indicating that inhomogeneous dynamical screening of electric fields by newly produced plasma is likely central to the observed emission [35,36]. While important details remain open, one aspect is robust: coherent radio emission from the inner magnetosphere requires sustained $e^\pm$ pair production [37].

In standard isolated radio pulsars, pair production is sustained via single-photon magnetic conversion, $\gamma + B \rightarrow e^- + e^+$, where the high-energy photon is produced as curvature radiation by locally accelerating charges. The optical depth τ is given by [38,39]

$$\tau \simeq 0.23 \alpha_{\text{EM}} m_e \int dx \sin \psi \left(\frac{B}{B_Q} \right) e^{-\frac{8 B_Q m_e}{3 \omega B \sin \psi}}, \quad (12)$$

where $B_Q \simeq 4.4 \times 10^{13}$ G, B is the magnetic field strength, ω the photon energy, and ψ the angle between the photon momentum and the magnetic field (with $\psi \sim \gamma^{-1} \ll 1$). The optical depth is exponentially suppressed unless ω is sufficiently large. Using the characteristic energy of curvature radiation $\langle \omega \rangle \sim 3\gamma^3/(2\rho_c)$, one finds that primary electrons must be accelerated to Lorentz factors of order $\gamma_T \sim 6 \times 10^7 \times (B_Q/B)^{1/3} (\rho_c/10^8 \text{ cm})^{1/3}$ to efficiently trigger pair production [40,41].

The maximal voltage drop along a field line for a NS with polar-cap radius r_{pc} is roughly [42]⁴

$$\Delta V \simeq \frac{B \Omega r_{\text{pc}}^2}{2}. \quad (13)$$

This corresponds to a maximal Lorentz factor

$$\gamma_{\text{max}} \sim \frac{\Delta V}{e m_e} \sim 10^7 \left(\frac{B}{10^{12} \text{ G}} \right) \left(\frac{1 \text{ s}}{P} \right)^2. \quad (14)$$

For rapidly rotating NSs, $\gamma_{\text{max}} \gg \gamma_T$, while as pulsars age and spin down, they can eventually reach a point where acceleration becomes too inefficient to ignite pair production. Since pair production is central to radio emission, pulsars become radio quiet and “die” once pair cascades can no longer be sustained.⁵

Axion hair modifies this picture by producing an additional potential drop that can be much larger than implied by Eq. (13) and, crucially, is not tied to the rotational frequency

³Here we fixed $\sin \psi \sim 10^{-3}$ and identified when the exponential suppression becomes ~ 0.1 [40].

⁴Pair production typically limits the realized voltage drop to be much smaller than this upper bound.

⁵We restrict attention to curvature-radiation-induced pair production relevant for rotation-powered pulsars. Magnetars and millisecond pulsars may instead be supported by inverse Compton or two-photon pair production [41,43] and are excluded from our analysis.

of the pulsar. As a result, sufficiently slowly rotating pulsars can sustain pair production with comparable efficiency to much younger objects. The nonobservation of such slowly rotating radio pulsars can therefore be used as evidence against the presence of long-range axion hair.

B. Identifying the death valley

The discussion above provides a qualitative understanding of pulsar death. Here we outline a quantitative calculation capturing the impact of axion-induced electric fields on the pulsar population.

For pulsars with active pair production, electric fields are largely screened and the current flow is set by the large-scale magnetic structure, $\vec{j} \simeq \nabla \times \vec{B}$, approaching a force-free configuration. Deviations arise when the supplied charges cannot simultaneously support the required current and screen the accelerating field: an electric-field gradient forms, and if the voltage drop is sufficiently large, a pair discharge occurs (temporarily screening the field). As pulsars age, pair production eventually ceases, driving the system toward an inactive state with vanishing current and no coherent radio emission. The key question is whether axion hair allows sustained discharge in old pulsars that would otherwise be below the conventional death line.

Using a one-dimensional approximation, one can combine the Lorentz force equation, Gauss’s law, and Ampère’s law (momentarily neglecting axion gradients and radiative losses) into a differential equation governing the evolution of the Lorentz factor along a field line [44],

$$\left(\frac{d\gamma}{ds} \right)^2 = 2 \left[\alpha_0 \sqrt{\gamma^2 - 1} - \gamma + 1 \right]. \quad (15)$$

Here s is the distance along the field line (normalized to the GJ skin depth), $\vec{j}_{\text{m}} \equiv \nabla \times \vec{B}$ is the current density fixed by the large-scale solution, and $\alpha_0 \equiv j_{\text{m}}/\rho_{\text{GJ}}$ is the discharge parameter. For $\alpha_0 \geq 1$, $E_{\parallel}$ grows with altitude and particles are accelerated until pair production occurs or geometric corrections suppress acceleration. For $0 \leq \alpha_0 < 1$, there exists a maximum Lorentz factor $\gamma^{\text{max}} = (1 + \alpha_0^2)/(1 - \alpha_0^2)$ beyond which particles cannot be accelerated and discharges do not occur. In force-free solutions, α_0 is a spatially varying $\mathcal{O}(1)$ function across open field lines; we give explicit examples and discuss field-line dependence in the companion paper [20].

Generalizing to include axion hair and curvature losses,⁶ one finds [45]

$$\frac{d^2\gamma}{ds^2} = \frac{j_{\text{m}}}{\rho_{\text{GJ}}} \frac{\gamma}{\sqrt{\gamma^2 - 1}} - 1 - \frac{\rho_a}{\rho_{\text{GJ}}} - \mathcal{R}_\gamma, \quad (16)$$

⁶We work with the second-order form rather than the square of the first-order equation.

where $\mathcal{R}_\gamma$ accounts for curvature-radiation losses [46,47],

$$\mathcal{R}_\gamma = \frac{8}{3} \frac{e^2}{m_e \rho_c^2} \gamma^3 \frac{d\gamma}{dt}, \quad (17)$$

with ρ_c the radius of curvature. For middle-aged and old pulsars, $|\rho_a/\rho_{\text{GJ}}| \gg 1$ near the star, motivating the intuition that the near-surface behavior resembles Eq. (15) with $\rho_{\text{GJ}} \rightarrow \rho_a$, yielding an effective discharge parameter $\alpha^{\text{eff}} = \pm j_m/\rho_a \sim 0^\pm$ (the sign depends on the magnetic pole). Because the axion profile is exponentially suppressed beyond $r \sim m_a^{-1}$, the system transitions back to the ordinary α_0 regime at larger radii, and new null surfaces defined by $\rho_{\text{GJ}} + \rho_a = 0$ can arise. The full discharge phenomenology across poles, including null-surface formation and return-current consistency, is established and validated with PIC simulations in the companion paper [20].

Let us summarize the qualitative outcomes relevant for pulsar demographics, focusing first on regions of the polar cap with $\alpha_0 > 0$ (return currents $\alpha_0 < 0$ are addressed below).

Without loss of generality take $\vec{\Omega} \cdot \vec{B} > 0$ so that $|\rho_{\text{GJ}}| < 0$:

- (i) Northern pole: $\rho_a > 0$. A null surface forms where $\rho_a = -\rho_{\text{GJ}}$. The current is carried by electrons, whereas screening ρ_a would require positive charges. As a result, an accelerating field develops and, for sufficiently large ρ_a , pair discharge is triggered. This persists even for pulsars that would be below the conventional death line.
- (ii) Southern pole: $\rho_a < 0$. The axion-induced effective charge density can be screened by slowly moving electrons, and beyond $r \gtrsim m_a^{-1}$ particles are accelerated primarily by the rotationally induced field. In this regime, long-range axion hair can effectively displace the location of the near-surface accelerating region and for sufficiently light axions can reduce the available voltage in one hemisphere.

To shift the death line in a globally self-consistent way, pair production must also remain active on field lines carrying return currents, $\alpha_0 < 0$; otherwise the global solution can collapse to an inactive disk-dome state. In simulations, gaps often open at altitudes $\mathcal{O}(R_{\text{NS}})$ and ignite once local conditions allow discharge. Distinguishing the two poles:

- (i) Northern pole: $\rho_a > 0$. The return-current region can supply the required current while screening the axion-induced field near the surface; particles are then accelerated mainly by the rotationally induced field at $r \gtrsim m_a^{-1}$, and discharge can fail for pulsars below the standard death line.
- (ii) Southern pole: $\rho_a < 0$. The rotationally induced and axion-induced fields can act constructively, enabling return-current particles to reach discharge threshold even when rotation alone would be insufficient.

Collectively, this implies that pulsars conventionally considered “dead” can remain radio active out to much larger periods, yielding the demographic tension emphasized

above. A notable consequence of the polarity asymmetry is a generic tension between long-range axion hair and the existence of interpulses in near-orthogonal rotators; PSR J1906 + 0746 [48,49] provides an important example, and the detailed interpulse analysis is presented in the companion paper [20].

To identify the corresponding “death valley,”⁷ we solve Eq. (16) over a range of pulsar parameters. At each step we compute the highest-energy curvature photon emitted by the primary particle, propagate it outward (including Schwarzschild effects for a $M = 1.4M_\odot$ star [50,51]), and evaluate the optical depth using Eq. (12). We define $\ell_{e^\pm}$ as the distance traveled by the photon when the optical depth reaches unity and define the pair-production distance as $d_{\text{pp}} \equiv \ell_{\text{acc}} + \ell_{\text{rad}} + \ell_{\text{mfp}}$, where ℓ_{acc} is the acceleration distance, ℓ_{mfp} satisfies $\tau(\ell_{\text{mfp}}) = 1$, and

$$\ell_{\text{rad}} \sim \sqrt{\frac{8\pi}{3}} \sqrt{\frac{\omega}{\omega_c}} \frac{\rho_c e^{\omega/\omega_c}}{e^2 \gamma}. \quad (18)$$

The gap height is determined by minimizing d_{pp} along the trajectory; pulsars with nonfinite gap heights are classified as dead. The full numerical implementation, as well as a simple analytical recipe, can be found in the Appendixes.

We map the results into the P - $\dot{P}$ plane assuming $\dot{P} \simeq \beta B^2/P$, with $\beta = \pi^2 R_{\text{NS}}^6/I_{\text{NS}} \simeq 6 \times 10^{-40} \text{ s/G}^2$ for canonical NS parameters. We compare in Fig. 1 death lines without (yellow) and with (red) axion hair for a light QCD axion with $\epsilon = 10^{-2}$ and $f_a = 10^{16} \text{ GeV}$. The figure highlights the main conclusion: axion hair induces a large, observationally incompatible shift in pulsar death.

IV. RESULTS AND CONCLUSIONS

The preceding section allows one to estimate, for a fixed set of pulsar parameters, whether a pulsar produces near-field radio emission. The remaining question is whether one can observationally distinguish a pulsar population whose lifetimes have been extended by axion hair from a population evolving under standard electrodynamics.

Figure 1 suggests that the observed pulsar population broadly aligns with standard expectations for pulsar death; however, one must confirm that the absence of high-period radio pulsars is not dominated by observational biases. To test this, we perform a population synthesis analysis (following, e.g., Refs. [65–70]) by: (i) sampling pulsar birth locations, (ii) drawing pulsar parameters at birth from parametrized distributions, (iii) determining present-day locations by sampling ages and evolving motion in the Galactic potential, (iv) evolving magnetorotational spin-down to obtain present-day P , $\dot{P}$ and magnetic fields,

⁷The term death valley, rather than death line, is often used since there is no unique mapping between $\dot{P}$ and the existence of $e^\pm$ production.




FIG. 2. Constraints on the light QCD axion (left), set at a confidence level of greater than 95% (see text for details), and the linearly coupled axion (right), set at the 95% confidence level. Gray regions are excluded by existing constraints [7,9,13,14,52–64]. Note that for the QCD axion we assumed an untuned axion-photon coupling $C \gtrsim 10^{-2}$. Note further that, for the linearly coupled axion, we take the strongest constraints in this mass range of the axion-photon coupling, which are set by a combination of pulsar polar caps [58], magnetic white dwarf polarization [59], NuSTAR observations of M82 and M87 [60], and Chandra observations of H1821 + 643 [61]. Red regions highlight the pulsar death analysis presented here. We truncate the death-line analysis at $m_a = 10^{-9}$ eV; for larger masses the axion-induced voltage drop becomes comparable to the transverse scale of the return-current field-line bundle and the global magnetospheric response may be altered (see discussion below). Green lines in the right panel show constraints (solid) from interpulses in PSR J1906 + 0746; the interpulse modeling and the sign dependence are presented in the companion paper [20]. The yellow line in the left panel corresponds to the conventional QCD axion relation.

(v) applying observational selection cuts (radio flux thresholds, beaming geometry, pulse width, survey sensitivity, etc.), (vi) comparing simulated and observed populations via a test statistic, and (vii) optimizing over free parameters. We test multiple distributions and assumptions to (1) show that minimal models can fit observations, (2) demonstrate that models without a death line are disfavored, and (3) assess the importance of each assumption. All technical details of the population synthesis, selection modeling, and statistical procedure are provided in the Appendixes.

To infer axion-compatible parameter space, we repeat the population synthesis using death lines specific to each model. We apply a conservative 50% visibility cut for pulsars below the line to account for hemispheric asymmetries in pair production. For each parameter set, we generate up to 10 mock populations, compute the mean and standard deviation of pulsars below the death line, and identify where a 2σ downward statistical fluctuation exceeds 14—corresponding to a 2σ excess of Australia Telescope National Facility (ATNF) pulsars below the fiducial line (limits on the light QCD axion are highly insensitive to this threshold). The resulting exclusions are shown in Fig. 2.⁸ Recall that the region of parameter space being probed for the light QCD axion is bounded to $f_a \leq 10^{17}$ GeV, $m_a \lesssim 10^{-9}$ eV, and $\epsilon \lesssim 0.07$ —within this range, all parameter space is excluded at greater than 95% confidence level (note

that this is not exactly at 95%, but greater than this value—this is because one cannot recover the conventional death line by smoothly varying, e.g., ϵ at fixed axion mass). Note that, for our fiducial analysis, we have chosen to not consider axion masses $m_a \gtrsim 10^{-9}$ eV. This is because our one-dimensional treatment is only valid when the transverse scale of the pair producing field lines can be neglected—for return currents, this implies pair production should occur at distances away from the surface $L_{pp} \gtrsim r_{pc}/\zeta$, with $\zeta \gtrsim 2$ characterizing the fractional size of the polar cap spanned by the return-current field lines. For low-mass axions, the potential drop is localized at distances larger than this threshold. For heavy axions, however, this is not the case. In order to determine where this characteristic threshold cut should be, we adopt the analytic formalism described in the Appendixes to compute death lines for a variety of axion masses, but inducing an artificial shift in the surface of the neutron star by a distance $r_{pc}/2$ —this effectively serves to remove the near-field axion potential drop from the analysis and allows us to assess only the impact of the exponentially suppressed tail.

Let us highlight that the constraints derived on the light QCD axion have been obtained using an axion-photon coupling $g_{a\gamma\gamma} = C\alpha_{EM}/(2\pi f_a)$, with $C = 1$. More generally, $C \equiv E/N - 1.92$, with E/N being the ratio of electromagnetic to color anomaly coefficients. For the KSVZ, DSVZ-I axions, this would imply $C \sim -1.92$ and $C \sim 0.75$ respectively. In general, $\mathcal{O}(1)$ variations in C induce very small shifts in the death lines, but have no impact on the derived constraints (this is because the region of the $P - \dot{P}$

⁸We have verified that constraints on the axion-photon interaction for the light QCD axion are subdominant to those presented here.

diagram driving the constraints is typically quite far from the derived death lines). By recomputing the death lines with suppressed values of C , and comparing with synthesized pulsar populations, we estimate that our constraints are expected to be valid for values of $C \gtrsim 10^{-2}$; for smaller values, which require a fine-tuned cancellation between the anomaly coefficients and the contribution from the pion, our constraints would not apply.

We have shown that axion gradients sourced from dense nuclear matter can decouple the electric field responsible for acceleration and pair production in the inner magnetosphere from the rotational period of the pulsar itself, leading to a population of slowly rotating radioactive stars and, in some cases, an incompatibility with interpulses. The nonobservation of such a population leads to competitive constraints on the light QCD axion and on linear axion-nucleon interactions, as summarized in Fig. 2. The companion paper [20] provides the complete electrodynamic derivations, PIC simulation details, and extended systematics underlying these results, while the Appendixes contain the population synthesis and statistical inference procedure and additionally draw on Refs. [71–102].

ACKNOWLEDGMENTS

We thank Anson Hook, Konstantin Springmann, and Kai Bartnick for enlightening discussions and Ani Prabhu for their useful comments on the manuscript. S. J. W. acknowledges support from a Royal Society University Research Fellowship (URF-R1-231065). This work is also supported by the Deutsche Forschungsgemeinschaft under Germany’s Excellence Strategy EXC 2121 Quantum Universe 390833306. A.C. is supported by an ERC STG grant (“AstroDarkLS,” Grant No. 101117510). A.C. acknowledges the Weizmann Institute of Science for hospitality at different stages of this project and the support from the Benozio Endowment Fund for the Advancement of Science. S.S. acknowledges financial support from the Spanish Ministry of Science and Innovation (MICINN) through the Spanish State Research Agency, under Severo Ochoa Centres of Excellence Programme 2025–2029 (CEX2024001442-S). This work is also part of the R&D&i project PID2023-146686NB-C31, funded by MICIU/AEI/10.13039/501100011033/ and by ERDF/EU. IFAE is partially funded by the CERCA program of the Generalitat de Catalunya. A. Ch. is supported by Martin A. and Helen Chooljian Member Fund and the Fund for Memberships in Natural Sciences. This article/publication is based upon work from COST Action COSMIC WISPerS CA21106, supported by COST (European Cooperation in Science and Technology). This work was supported by a grant from the Simons Foundation (MP-SCMPS-00001470), as well as an Alfred P. Sloan Fellowship and a Packard Foundation Fellowship in Science and Engineering to A. P. This work was supported by the U.S. Department of Energy (DOE), Office of Science, National Quantum Information

Science Research Centers, Superconducting Quantum Materials and Systems Center (SQMS) under Contract No. DE-AC02-07CH11359. S.R. is supported in part by the U.S. National Science Foundation (NSF) under Grant No. PHY-2412361. The work of S. R. was also supported by the Simons Investigator Award No. 827042.

DATA AVAILABILITY

The data that support the findings of this article are not publicly available upon publication because it is not technically feasible and/or the cost of preparing, depositing, and hosting the data would be prohibitive within the terms of this research project. The data are available from the authors upon reasonable request.

APPENDIX A: SPECULATIONS ABOUT THE VANILLA QCD AXION

It remains an open question as to whether the vanilla QCD axion is sourced within the very high density matter in the core of the neutron star. If this is the case, probes of the neutron star equation of state may be one of the most promising ways to directly probe the QCD axion. However, the modification of the structure of neutron stars predominantly is confined to the core, making surface probes like cooling insensitive to sourcing of the QCD axion deep inside the star. Since the probe discussed here is sensitive to the axion gradient at the surface, and the QCD axion would be sourced by the core of the neutron star, it is unclear whether there would be sufficient leakage outside of the core to shift the pulsar death line. In order to directly




FIG. 3. Shift in the death line that would arise should the core of the neutron star be sufficiently dense to support the sourcing of the QCD axion (i.e., $\epsilon = 1$). Here, we have assumed a core radius of $R_c = 5$ km.

address this question, we assume the QCD axion is sourced within a core of radius $R_c \sim 5$ km and adopt a radial axion profile consistent with Eq. (6), but taking $R \rightarrow R_c = 5$ km. The shift in the death line produced from this profile is shown for two values of f_a in Fig. 3. Here, one can see that a fraction of parameter space near $f_a \sim 10^{17}$ GeV may still be accessible, despite the strong suppression of the axion field at the surface of the star. It is worth highlighting that we have chosen to truncate our analysis conservatively at $f_a \lesssim 10^{17}$ GeV, however, should QCD axions be sourced with slightly higher decay constants, these axions will leave discernible imprints on the pulsar population.

APPENDIX B: THE ONE-SIDED LIGHTHOUSE

One of the interesting features that is produced by these axion field gradients is an asymmetry in the sign of the axion-induced electric field across each of the poles. In particular, for the case of sufficiently light axion hair around active pulsars, screening can effectively turn off pair production along a large part of one pole, while leaving it unchanged in the other. A natural question to ask is whether there exist clear observational consequences that could arise from this asymmetry.

One possible avenue is to exploit double-pole pulsars—a rare class of nearly orthogonal rotating pulsars that allow for radio emission to be observed from both polar caps [48,49]. Depending on the relative alignment, the line of sight of Earth may pass through one pole where pair production is enhanced and another where it is suppressed. In the event that the line of sight passes both return currents, or both outflowing currents, one is guaranteed to observe at least one pole where pair production is screened (this follows from the fact that, for fixed ρ_{GJ} and α_0 , the axion charge density has an asymmetry). For nearly maximally misaligned rotators, there is also a possibility that one observes radio emission from outflowing current on one pole and return current on another; depending on the axion and field configuration (and which pole contains the outflowing and return currents), one may expect axion hair to either screen both emission zones or neither. Thus, the geometry of the magnetic field and the sign of the axion hair must be uniquely identified in order to derive constraints.

There exists a remarkable pulsar, PSR J1906 + 0746, for which interpulses are observed and which is in binary that induces a small precession such that different parts of the polar cap are observed over the timescales of years. The specific geometry of this configuration with respect to two polar caps is clearly illustrated in Fig. 1 of [71], where one can see the line of sight is typically⁹ traversing two return-current regions, both with the same sign of ρ_{GJ} , but with opposite signs of ρ_a . For sufficiently small axion

masses and sufficiently large values of ρ_a , one can ensure that axion hair suppresses pair production on one of these return currents. In addition, depending on the sign of ρ_a , one may also expect the precession to allow one to test the screening induced by axion hair along field lines with outflowing currents.

Let us derive an approximate condition for the axion mass and coupling at which this effect should arise, focusing momentarily on the effect on the discharge in the return-current zone (since these lead to an unambiguous signal, irrespective of the sign of ρ_a). PSR J1906 + 0746 has a rotational period of ~ 0.144 s, a spin-down rate of $\dot{P} = 2.0 \times 10^{-14}$, an inferred surface magnetic field strength of 1.7×10^{12} G, and an inclination of $\alpha = 99.41^\circ$ [49]. For a rotational period $P \sim 0.144$ s, the conventional death line is sitting near $\dot{P} \sim 10^{-18}$; since $\dot{P} \propto B^2$, one can infer that an identical pulsar with a magnetic field $B \lesssim 10^{10}$ G would be inactive. For dipolar fields, the magnetic field scales as $B \propto (R_{\text{NS}}/r)^3$, implying that the mean-free path of curvature photons in PSR J1906 + 0746 becomes finite only at a distance of ~ 4 stellar radii above the surface of the star. This implies that axion hair will suppress emission along one pole in PSR J1906 + 0746 if $|\rho_a(r)/\rho_{\text{GJ}}(r)| \gg 1$ for $r \lesssim 5R_{\text{NS}}$. For the light QCD axion, and assuming $g_{a\gamma\gamma} \sim \alpha_{\text{EM}}/(2\pi f_a)$, this amounts to the condition $m_a \gtrsim 8\pi \cos \alpha / (\alpha_{\text{EM}} P)$, which cannot be satisfied while ensuring $m_a \times 4 \times R_{\text{NS}} \ll 1$ (this statement, however, is not necessarily true for other pulsars). For the CP -violating axion, one finds instead in the low-mass limit

$$m_a \lesssim 10^{-12} \text{ eV}, \quad (\text{B1})$$

$$g_N \times g_{a\gamma} \gtrsim 8 \times 10^{-39} \text{ GeV}^{-1}. \quad (\text{B2})$$

This constraint is only marginally weaker than the constraint derived using the death line and is included (computing the constraint using the full field profile across all axion masses, not just in the massless limit) in Fig. 2 for comparison (solid line).

It is worth highlighting that this constraint becomes significantly stronger if axions also turn off the discharge along the field lines with outflowing currents. This is because acceleration necessarily takes place close to the surface and is limited by the traverse size of the polar-cap region; here, the condition that axion screening is sufficient to suppress emission is merely that ρ_a is large on distances scales $\gg 2r_{\text{pc}}$ from the stellar surface. Repeating the procedure above, but imposing $\rho_a(r)/\rho_{\text{GJ}}(r) \gg 1$ for $r \lesssim R_{\text{NS}} + 2r_{\text{pc}}$ instead leads to a sensitivity in the massless axion limit at the level of

$$m_a \lesssim 10^{-12} \text{ eV}, \quad (\text{B3})$$

$$g_N \times g_{a\gamma} \gtrsim 2 \times 10^{-39} \text{ GeV}^{-1}. \quad (\text{B4})$$

⁹Precision of the orbit causes the field-line bundle being probed to shift.

There is a subtlety, however; in order to adopt this stronger constraint, one must confidently probe outflowing field lines with $\text{sign}(\rho_a) = \text{sign}(\rho_{\text{GJ}})$. Observationally, there is no clear observational signature which allows one to differentiate the sign of ρ_{GJ} ; moreover, both $g_{a\gamma\gamma}$ and g_N can take either sign, meaning ρ_a is also undetermined. Therefore, the most promising route is to hope for the detection of a similar object that exclusively probes the outflowing field lines on both poles. In order to merely demonstrate the potential power of such a probe, we include a dashed line in Fig. 2 to demarcate the region of parameter space for which efficient screening along the outflowing field line should alter the emission (this is done for both models).

APPENDIX C: FIELD PROFILES FOR THE LINEAR INTERACTION CASE

Here we provide the fitting formulas for the axion profile in the linear model, analogous to Eq. (6). For $m_a \lesssim 10^{-11}$ eV the profile is

$$a \simeq 1.5 \times 10^{13} \text{ GeV} \left(\frac{g_N}{10^{-23}} \right) \frac{R_{\text{star}}}{r} \exp[-m_a(r - R_{\text{star}})]. \quad (\text{C1})$$

In the opposite limit, when the axion mass is large, $1/m_a \lesssim \text{km}$, a better fit to the numeric solution outside the star is

$$a \simeq 10^{11} \text{ GeV} \left(\frac{g_N}{10^{-23}} \right) \left(\frac{2 \times 10^{-10} \text{ eV}}{m_a} \right)^2 \times \frac{R_c}{r} \exp[-m_a(r - R_c)], \quad (\text{C2})$$

with $R_c \simeq 10.5$ km. We numerically verified this fit for masses up to 10^{-9} eV, where the field value outside the star is already quite small. For such heavy axion masses, the axion field profile has non-negligible dependence on the density profile of the star, in particular on the density within a distance of $1/m_a$ from the edge of the star. The axion profile above was found adopting the Akmal-Pandharipande-Ravenhall (APR) equation of state [72] and a NS with mass $1M_\odot$. When the inverse of the scalar mass is comparable to the size of the entire star, that is to say $m_a \sim \mathcal{O}(\text{few})10^{-11} - 10^{-10}$ eV, then the solution outside the star is well approximated by

$$a \simeq 2.5 \times 10^{11} \text{ GeV} \left(\frac{g_N}{10^{-23}} \right) \left(\frac{10^{-10} \text{ eV}}{m_a} \right)^2 \times \frac{R_{\text{star}}}{r} \exp[-m_a(r - R_{\text{star}})], \quad (\text{C3})$$

where $R_{\text{star}} = 12$ km and the numerical prefactor have been fixed using the APR density profile. Details about the

derivation of these formulas are provided in the companion paper [20].

In the following we provide details about the construction of the pulsar death valley, both numerically and with a simple analytical recipe, as well as our neutron star population modeling and the statistical analysis on which our bounds are based.

APPENDIX D: CONSTRUCTING THE PULSAR DEATH VALLEY

In this section we describe in detail how, for a fixed set of pulsar parameters, we determine if the pulsar is active. We start with our full, numerical solution and then we also provide an easy analytical solution which agrees extremely well with the numerics for the parameters of interest.

1. Full numerical solution

We start by fixing the surface magnetic field strength B_0 , the rotational frequency Ω , and the curvature of the field lines ρ_c of a given neutron star. Our fiducial analysis adopts a radius of curvature consistent with a dipolar field configuration,

$$\rho_c \simeq 9 \times 10^7 \left(\frac{\theta_0}{\theta} \right) \sqrt{\frac{P}{1 \text{ sec}}} \text{ cm}, \quad (\text{D1})$$

θ_0 is the angular opening of the polar cap and θ defines the field line of interest (note that we also run analyses fixing the value of ρ_c and varying this fixed value across a wide range of characteristic scales consistent with both dipolar and quadrupolar fields). We then fix the current demanded by the twist on the local field-line bundle under consideration $j_m \equiv |\nabla \times \vec{B}| = \alpha_0 \rho_{\text{GJ}}$; we adopt a prefactor $\alpha_0 = 2$ which is roughly consistent with characteristic values obtained in the force-free simulations of inclined pulsars [73].

We then evolve Eq. (16) to determine the acceleration of the primary particles along the field lines. We stop the acceleration at a distance scale $r \sim 2r_{\text{pc}}$, as this is where the one-dimensional approximation breaks down [74]. At each time step, one can compute the energy radiated per unit frequency per unit distance, which is given by [33]

$$dI = \frac{\sqrt{3}}{2\pi} \frac{e^2}{\rho_c} \gamma F(\omega/\omega_c) d\omega d\ell \quad (\text{D2})$$

where

$$F(x) \equiv x \int_x^\infty K_{5/3}(x) dx, \quad (\text{D3})$$

and $K_{5/3}$ is the modified Bessel function of the second kind. The characteristic photon energy emitted is $\omega_c \sim 3\gamma^3/(2\rho_c)$, but pair cascades are initiated by the highest-energy photons,

which are typically a factor of a few higher than this value. Taking the limit $\omega \gg \omega_c$, we can estimate

$$F(x) \sim \sqrt{\frac{\pi x}{2}} e^{-x} \left(1 + \frac{55}{72} \frac{1}{x} + \dots \right). \quad (\text{D4})$$

We only expect a γ -ray of energy ω to be radiated over a distance scale ℓ_{rad} when the total energy radiated exceeds the photon energy, i.e.,

$$\omega \leq \int_{\ell_0}^{\ell_0 + \ell_{\text{rad}}} \int_{\omega}^{\infty} dI(\omega') d\omega' d\ell. \quad (\text{D5})$$

Assuming that the radiation length is sufficiently small that γ can be treated as a constant, one can define an implicit relation between ℓ_{rad} and the maximum photon energy radiated given in Eq. (18). Notice that as $\omega/\omega_c \rightarrow \infty$ the radiation length $\ell_{\text{rad}} \rightarrow \infty$, but the mean-free path of the photon $\ell_{e^\pm}$, obtained by inverting the optical depth given in Eq. (12) goes to zero. Conversely, in the limit $\omega/\omega_c \rightarrow 1$, the radiation length shrinks, but $\ell_{e^\pm}$ becomes large. This implies an implicit trade-off between production of the photon and pair production, as one might expect.

At each point in the evolution of the primary particle, we compute the total distance traveled, and then project onto a two-dimensional location of a specific field line of interest. A photon of fixed energy ω is emitted from that location parallel to the magnetic field line itself, and the photon trajectory is traced by solving

$$\frac{dx^\mu}{d\lambda} = \frac{\partial \mathcal{H}}{\partial k_\mu}, \quad (\text{D6})$$

$$\frac{dk_\mu}{d\lambda} = -\frac{\partial \mathcal{H}}{\partial x^\mu}, \quad (\text{D7})$$

with $\mathcal{H} = g^{\mu\nu} k_\mu k_\nu$. Here, we adopt a Schwarzschild metric for a $M = 1.4M_\odot$ star with a radius of 12 km. For the sake of illustration, we also consider in an additional variant considered here the case of an aligned quadrupolar field—the procedure here is the same, however, the structure and strength of the magnetic field have been altered accordingly. Note that, for a purely quadrupolar field, the magnetic field lines follow the solution

$$\frac{dr}{r} = \frac{d\theta(3 \cos^2 \theta - 1)}{\sin 2\theta}, \quad (\text{D8})$$

which implies that open field lines will be confined to $\theta \lesssim 0.62 \times (R_{\text{NS}} \times \Omega)$ [75]. Since the open field bundle is notably reduced with respect to a dipolar pulsar, this tends to imply a larger characteristic radius of curvature of the field line. For this reason, we take $\rho_c \sim 10^9$ cm in our fiducial quadrupole analysis (although quadrupolar fields

can also have characteristic values of ρ_c much smaller than dipolar fields in certain regimes).

Since the gap height is effectively the sum of the acceleration length ℓ_{acc} , the radiation length, and the mean-free path, one can ascertain the energy of pair-producing photons and the explicit gap height by minimizing the sum of these three quantities over the path of the primary particle. The transition between a small but finite gap height and an infinite gap height (implying pair production never occurs) is extremely rapid, and thus in practice we define pulsar death as the point where the gap height, computed using the procedure outlined above, exceeds 2 km.

One can gain a bit of analytic intuition for this problem by noting that optical depth is exponentially suppressed until $\chi \equiv (\omega/m_e) \times (B/B_q) \times \sin \psi/2 \gtrsim \mathcal{O}(1/7)$, and ψ always remains small near the neutron star; collectively, these conditions suggests that if pair production occurs, the mean-free path should be approximately given by

$$\ell_{e^\pm} \sim \frac{2\rho_b}{7} \left(\frac{m_e}{\omega} \right) \left(\frac{B_q}{B} \right). \quad (\text{D9})$$

For a fixed value of γ , Eqs. (18) and (D9) suggest that the minimum of $(\ell_{\text{rad}} + \ell_{e^\pm})$ is roughly obtained by solving

$$\begin{aligned} & \left(\frac{\omega}{\omega_c} \right)^{5/2} e^{\omega/\omega_c} \left(1 - \frac{55}{72} \frac{\omega_c}{\omega} + \dots \right) \\ & \simeq 40 \times \left(\frac{\rho_c}{10^7 \text{ cm}} \right) \left(\frac{10^{12} \text{ G}}{B} \right) \left(\frac{10^7}{\gamma} \right)^2. \end{aligned} \quad (\text{D10})$$

For pulsars near the death line, this typically yields an energy $\omega \sim \mathcal{O}(5\text{--}10) \times \omega_c$. One can then use the optical depth to turn this into a condition on the boost factor of the primary particle. This estimate works reasonably well across much of the relevant parameter space, but does fail in certain regimes.

The discussion outlined above assumes that the primary pair-production channel is via single-photon $e^\pm$ pair production, with the incident high-energy γ ray being produced by curvature radiation. This is a valid assumption across a wide range of the pulsar parameter space (including for standard radio pulsars), however, it is not necessarily valid for certain subpopulations. Below, we briefly discuss two alternative pair-production channels.

For high magnetic field pulsars and magnetars, resonant inverse Compton scattering (often referred to as “RICS”) can provide an alternative plasma loading mechanism [41,76–78]. Here, hot thermal photons emitted from the stellar surface (in particular, photons emitted from the so-called “hot spots,” which are generated as highly energetic particles, are accelerated into the surface of the star) can up-scatter off of primary charges accelerated from the stellar surface or off of secondary plasma particles initiated by curvature radiation—the final state photon sits at pair-producing energies and will thus contribute to the




FIG. 4. Left: pulsar death line computed using a dipolar field and a radius of curvature fixed to the last open field line (yellow), a dipolar field, but imposing a radius of curvature of $\rho_c = 10^6, 10^7$, or 10^8 cm (green; solid, dashed, and dot-dashed), or an order-one quadrupole-dipolar field with a fixed radius of curvature of $\rho_c = 10^9$ cm (blue). Right: pulsar death line in axion models. Red curve (bold) shows the light QCD axion with $\epsilon = 10^{-2}$ and $f_a = 10^{14}$ GeV, while the green curve shows the death line when $m_a \gg R_{\text{NS}}$, corresponding to the limit when $\epsilon \rightarrow 0$. In the case of the linear axion-nucleon coupling (labeled “linear int.”), we fix the axion mass to $m_a = 10^{-12}$ eV and vary the product of the couplings $g_{a\gamma} \times g_N$ (see small number next to each line, shown in units of GeV^{-1}). In all cases the radius of curvature is fixed to that of the last open field line.

generation of plasma via the same single-photon pair-production process discussed above. Resonant scattering off of secondary plasma can play an important role in establishing the maximum multiplicity of pair cascades in strong field pulsars [41], but is less interesting from the prospective of pulsar death, as pair cascades must be initiated from curvature photons. Resonant scattering off primary particles, on the other hand, typically have a multiplicity that is too low to efficiently screen the electric field [41,78], implying they are not relevant from the perspective of establishing pulsar death.

In Fig. 4, we demonstrate how the death line shifts in response to changes in the underlying assumptions. In particular, the left panel of Fig. 4 shows the impact of fixing the radius of curvature to $\rho_c = 10^6$ cm (green, solid), $\rho_c = 10^7$ cm (green, dashed), and $\rho_c = 10^8$ cm (green, dot-dashed), or including a quadrupole moment aligned with the dipole, and with a characteristic field strength equal to the dipolar field strength at the surface. Since a quadrupole inherently alters the curvature and structure of the open field-line bundle, we fix the radius of curvature in this example to $\rho_c = 10^9$ cm (blue).

In the right-hand panel of Fig. 4, we show a variety of different death lines computed for both the light QCD axion and the linearly coupled axion. For the light QCD axion, we note that the effect of the axion gradient has two terms

$$g_{a\gamma} \partial_r a \sim \frac{\alpha_0}{2R_{\text{NS}}} (1 + m_a R_{\text{NS}}). \quad (\text{D11})$$

For small axion masses (or equivalently when $\epsilon \rightarrow 0$), the axion-induced voltage drop tends toward a fixed value, while for large axion masses the effect scales proportionally to the mass itself. Since there exists a maximal value of m_a for a fixed f_a , we expect the death lines for the light QCD axion to asymptote to a fixed value when $m_a \times R_{\text{NS}} \ll 1$ and be maximally deviating when $\epsilon \sim 0.07$ (the maximal value allowed in our analysis). Therefore, we only plot these two limits.

In the case of the linearly coupled axion (blue lines), we instead fix the axion mass and vary the product of the axion-nucleon and axion-photon coupling. Here, we can see the shift in the death line scales directly with this product, as one might expect.

2. A simple analytical solution

Before delving into the matter of neutron star populations and pulsar data, we find it useful to provide also a simple, analytical solution for the problem of pair discharge and the determination of a pulsar death line.

As we reported in the main text, the equation to be solved reads

$$\frac{d^2\gamma}{ds^2} = \frac{e}{m_e} \left(\frac{\gamma \alpha_0 \rho_{\text{GJ}}}{\sqrt{\gamma^2 - 1}} - \rho_{\text{GJ}} - g_{a\gamma} B \partial_r a \right) - \mathcal{R}_\gamma, \quad (\text{D12})$$

where s is the distance from the NS surface, γ is the gamma factor for the accelerated electrons, and $\mathcal{R}_\gamma$ accounts for radiative losses, see Eq. (17). The boundary condition is that the electrons start at rest at $s = 0$. We start considering the case *without* axions ($g_{a\gamma} = 0$), fix $\alpha_0 = 2$, and approximate $\gamma/\sqrt{\gamma^2 - 1} \approx 1$, which is a good approximation given that charges will be accelerated to large values of γ pretty quickly. Hence the equation simplifies to

$$\frac{d^2\gamma}{ds^2} = \frac{e}{m_e} \rho_{\text{GJ}} - \mathcal{R}_\gamma. \quad (\text{D13})$$

One can then identify two different regimes: the regime where the GJ charge density *always* dominates over the curvature radiation backreaction and the electrons keep accelerating up to $2r_{\text{cap}}$; the regime where the backreaction is large enough that at some point $\frac{d^2\gamma}{ds^2} \sim 0$ and therefore $d\gamma_e/ds \sim \text{const}$ (this corresponds to the radiation-reaction limited regime, mostly expected in millisecond pulsars). Here we consider only the first regime, which, as we will show, describes well a pulsar with large periods. The equation becomes

$$\frac{d^2\gamma}{ds^2} = \frac{e \rho_{\text{GJ}}}{m_e}, \quad (\text{D14})$$

whose solution yields an analytical estimate for the maximum Lorentz factor at $2r_{\text{pc}}$,

$$\gamma_{\text{max}} \simeq \frac{8\pi e}{P m_e} B_{\text{ref}} \left(\frac{r_{\text{pc}}^2}{1 + 2r_{\text{pc}}/R_{\text{NS}}} \right) \sqrt{\frac{P}{s} \left(\frac{\dot{P}}{10^{-15}} \right)}, \quad (\text{D15})$$

where we used the standard scaling

$$B = B_{\text{ref}} \left(\frac{P}{s} \frac{\dot{P}}{10^{-15}} \right)^{1/2}. \quad (\text{D16})$$

Then, we can impose an approximate condition for pair production as

$$\chi \equiv \frac{3 \xi_\gamma \gamma_{\text{max}}^3 s_{\text{min}} B(s_{\text{min}})/\rho_c}{2m_e \rho_c B_Q} \simeq \chi_{\text{min}} \simeq 1/7, \quad (\text{D17})$$

where $\chi_{\text{min}} \simeq 1/7$ has been tuned to the numerical solution of $\tau = 1$, ξ_γ is taken to be an order-one coefficient which accounts for the correction between the pair-producing photon energy and the typical photon energy ($\omega \equiv \xi_\gamma \omega_c$), and where s_{min} is the maximum of the function $sB(s)$. For a dipolar magnetic field this reads $s_{\text{min}} = R_{\text{NS}}/2$. Plugging Eq. (D15) into (D17) we find the simple scaling

$$\begin{aligned} \dot{P} &\sim 8 \times 10^{-18} \left(\frac{P}{s} \right)^{5/2} \left(\frac{1.3 \times 10^{12} \text{ G}}{B_{\text{ref}}} \right)^2 \\ &\times \left(\frac{\chi_{\text{min}}}{1/7} \frac{5}{\xi_\gamma} \right)^{1/2} \left(\frac{12 \text{ km}}{R_{\text{NS}}} \right)^5, \end{aligned} \quad (\text{D18})$$

which we checked to reproduce very well the numerical results for $P \gtrsim 0.5$ s.

Let us now add the axions into the picture, starting with the case of very small axion masses. In the region of parameter space where axions are sourced, the axion effective charge density always dominates over the GJ charge density; therefore, the differential equation to solve is simply

$$\frac{d^2\gamma_e}{ds^2} = -\frac{e g_{a\gamma}}{m_e} B(s) \partial_r a \simeq \frac{e B(s)}{2m_e} \frac{\alpha_{\text{EM}} R_{\text{NS}}}{(R_{\text{NS}} + s)^2}, \quad (\text{D19})$$

where we neglected the exponential factor in the axion field (which is approximately 1 for small axion masses), and we used $g_{a\gamma} \sim \frac{\alpha_{\text{EM}}}{2\pi f_a}$. Using the same procedure and reference values as in the case without axions, we find

$$\begin{aligned} \dot{P} &= 4 \times 10^{-19} \left(\frac{P}{s} \right) \left(\frac{1.3 \times 10^{12} \text{ G}}{B_{\text{ref}}} \right)^2 \\ &\times \left(\frac{\chi_{\text{min}}}{1/7} \frac{5}{\xi_\gamma} \right)^{1/2} \left(\frac{12 \text{ km}}{R_{\text{NS}}} \right)^{7/2}, \end{aligned} \quad (\text{D20})$$

which again agrees very well with the full numerical results for $P \gtrsim 1$ s. One can also do the same computation for the massive axion case, solving the simplified equation (valid for $m_a R_{\text{NS}} \gg 1$)

$$\frac{d^2\gamma_e}{ds^2} \simeq \frac{\alpha_{\text{EM}} e B(s)}{2m_e} e^{-m_a s} m_a. \quad (\text{D21})$$

In this case we find

$$\begin{aligned} \dot{P} &= 8 \times 10^{-22} \left(\frac{P}{s} \right)^{1/4} \left(\frac{1.3 \times 10^{12} \text{ G}}{B_{\text{ref}}} \right)^2 \\ &\times \left(\frac{\chi_{\text{min}}}{1/7} \frac{5}{\xi_\gamma} \right)^{1/2} \left(\frac{12 \text{ km}}{R_{\text{NS}}} \right)^{11/4}. \end{aligned} \quad (\text{D22})$$

In all three cases (without axions and with either massless or massive axions), the analytical formulas reproduce the numerical results extremely well, provided that radiation losses are negligible. This condition holds for large periods, i.e., on the right-hand side of the “kink” feature in the full numerical curves.

For smaller periods, the friction term during the acceleration cannot be neglected, and in general it becomes more difficult to derive simple analytical results. However, in the massive axion case, owing to the abrupt fall of the axion profile and the associated acceleration, one can still obtain a




FIG. 5. $P - \dot{P}$ diagram comparing the numeric results for the standard fiducial analysis (yellow) and light QCD axion in the small (green) and large mass limit (red) with our analytic approximations (dashed lines). For the small mass limit of the light QCD axion we only display the large P derivation, as the transition observed below $P \lesssim 1$ s has multiple competing scales which make analytic estimates less tractable.

reasonably good analytical approximation. To this end, we solve Eq. (D21) and the equation of motion neglecting friction, up to the point at which friction becomes dominant and γ_e saturates. To estimate where this transition occurs, we adopt a simple energetic argument, imposing

$$\int_0^{s_1} ds \left(\frac{\alpha_{\text{EM}} e B(s)}{2m_e} e^{-m_a s} m_a \right) = \int_0^{s_1} ds \mathcal{R}_\gamma, \quad (\text{D23})$$

where $\mathcal{R}_\gamma$ is evaluated on the solution of Eq. (D21). Following this procedure we get

$$\dot{P} = 6 \times 10^{-21} \left(\frac{P}{s} \right)^{-5/7} \left(\frac{\rho_c}{9 \times 10^7 \text{ cm}} \right)^{4/7} \left(\frac{1.3 \times 10^{12} \text{ G}}{B_{\text{ref}}} \right)^2 \times \left(\frac{\chi_{\text{min}}}{1/7 \xi_\gamma} \right)^{8/7} \left(\frac{12 \text{ km}}{R_{\text{NS}}} \right)^{8/7}. \quad (\text{D24})$$

In Fig. 5 we compare the analytic approximations obtained in Eqs. (D18), (D20), (D22), and (D24), with the numerical results outlined in the preceding sections. We show results for the standard death line, computed without an axion (yellow), and two death lines for the light QCD axion, one obtained in the low-mass $m_a R_{\text{NS}} \ll 1$ (green) limit and one obtained in the high-mass limit $m_a R_{\text{NS}} \gg 1$ (red). Analytic estimates are shown in dashed lines. In both the standard scenario and the massless axion limit, analytic




FIG. 6. $P - \dot{P}$ diagram showing analytic derivation of death line produced under the assumption that the axion-induced voltage drop on distances $r \leq r_{\text{pc}}/2$ is zero; this analysis is performed in order to ensure that the one-dimensional treatment of pair production on the return current remains valid, and justifies the cutoff scale of $m_a \leq 10^{-9}$ eV. The neutron star population corresponds to one mock population realization of model 10 (which is the best-fitting population that does not impose a death line cutoff).

approximations are only applicable above the break in the power law, while for the heavy axion we plot the limiting approximations at $P \rightarrow \infty$ and $P \rightarrow 0$. All analytic results appear to be in excellent agreement with numerical calculations.

As mentioned in the proceeding section, we impose a cutoff threshold on the axion mass, requiring $m_a \leq 10^{-9}$ eV. This comes from the fact that the longitudinal extent of the axion-induced voltage drop becomes smaller than the cross-sectional field-line bundle of the return current. In order to assess this cutoff threshold more precisely, we use the analytic approach here to rederive approximate death lines under the assumption that the axion induces no additional electric field on distances $r \leq r_{\text{pc}}/2$ from the neutron star surface. We plot these revised death lines in Fig. 6 for various axion masses. Instead of showing the ATNF pulsar population, we plot instead one mock population realization produced using model 10 of Appendix F, which corresponds to the best-fitting model that does include a hard cutoff in the death line. At low axion masses ($m_a \lesssim 10^{-10}$ eV), the new cutoff has no effect, while at higher axion masses the death line weakens for lower-period pulsars. The adopted threshold cut of $m_a \leq 10^{-9}$ eV can be seen by eye to be conservative, as it is still in strong tension with nearly all of the pulsars sitting below the standard death line.




FIG. 7. Proxy for the pulsar age $\tau_{\text{age}} \sim P/(2\dot{P})$ (left), flux density (center), and inferred pulsar luminosity (right) at $\nu = 1.4$ GHz as a function of effective distance from the death line $X_{\text{death}} = \pm \text{Min}[\sqrt{(\log P - \log P_{\text{death}})^2 + (\log \dot{P} - \log \dot{P}_{\text{death}})^2}]$, where the $+/-$ corresponds to pulsar above/below the death line (i.e., the death valley has been highlighted in gray), and we have used the fiducial death line shown in Fig. 1. Magnetars and millisecond pulsars have been removed from the dataset, and the Crab pulsar has been highlighted with a red star to serve as a reference value. In the right panel, a power law has been fit to the data to illustrate the general trend.

APPENDIX E: IS THERE EVIDENCE FOR A DEATH LINE?

In this appendix the question we seek to address is whether there is clear evidence for a death line in the pulsar distribution (at least in some rough definition). Only after having established the need for such a line, one may then study if the presence of axion hair unambiguously shifts where such a death line would lie.

Let us recall that, from a theoretical perspective, the existence of a pulsar death line (or, more appropriately, a death valley) is apparent: pair cascades are a crucial ingredient in the generation of near-field coherent radio emission, and the generation of pair cascades requires a sufficient number density of very high-energy radiation to be present near the star, something which can only be sustained for pulsars with sufficiently high magnetic fields and rotational frequencies. Pulsars lose their rotational energy and their magnetic field strength as they age, and thus one expects old pulsars to become radio quiet. This correlation can be seen by noting that a rough proxy for the pulsars age is given by $\tau_{\text{age}} \sim P/(2\dot{P})$, which follows from integrating $\dot{P}$ assuming negligible initial spin period and a constant value of B , and plotting τ_{age} as a function of some distance modulus $X_{\text{death}} = \pm \text{Min}[\sqrt{(\log P - \log P_{\text{death}})^2 + (\log \dot{P} - \log \dot{P}_{\text{death}})^2}]$ which allows one to define some notion of how close a pulsar is to the death line. Here, the $+/-$ sign corresponds to pulsars above/below the death line, and we have used the fiducial death line shown in Fig. 1. A scatter plot for the pulsar distribution (neglecting millisecond pulsars (MSPs) and magnetars) is shown in the left panel of Fig. 7. Now, assuming the birth rate of pulsars at times $t \gtrsim \mathcal{O}(100)$ Myr is not significantly lower than it is today (which would be extremely odd, given that the star formation rate was much more active in the past), one would expect more than an

order of magnitude more pulsars with ages above 100 Myr than with less than 100 Myr. Figure 7, however, shows that this population is minuscule, suggesting that the old pulsar population is simply not observed.

Now the next question that immediately arises is whether the lack of observation of the old pulsar population could simply be a limit of current telescopes observational sensitivity—or said equivalently, are old pulsars simply too faint to be observed? The short answer is that there exists no observational evidence to suggest that this is the case. If one plots the observed radio flux density (at, say, $\nu = 1.4$ GHz) as a function of the death distance modulus X_{death} , one sees only a very mild correlation, with the general scaling only a small effect relative to the scatter in the population—this is shown explicitly in the central panel of Fig. 7.¹⁰ At this point, one may try to argue that older pulsars should be preferentially close to Earth, and perhaps the choice of plotting the flux density rather than the intrinsic luminosity has somehow disguised a more prominent feature that arises near the death line. The radio luminosity for each pulsar can be inferred by inverting the equation for the flux density

$$S \simeq \frac{L}{d^2 d\Omega}, \quad (\text{E1})$$

where d is the inferred distance of each pulsar and $d\Omega = 4\pi(1 - \cos \rho_b)$ is the fractional sky coverage of the radio beams. Here, ρ_b is the angular opening of each beam, which from geometric considerations is roughly given by

¹⁰When flux density measurements at this frequency are not available, we extrapolate from observations at 400 MHz or 2 GHz using the typical frequency scaling relation $S \propto \nu^{-1.4}$ to provide an approximate estimate. This rescaling has no impact on the inferred trend in this figure.

$$\rho_b \simeq \sqrt{\frac{9\pi h_{\text{em}}}{2P}}, \quad (\text{E2})$$

with h_{em} being so-called “emission height” (the emission height being a proxy for the height above the surface at which the radio emission last scattered¹¹). For typical pulsars, the inferred emission heights tends to take on values $h_{\text{em}} \sim 300$ km, with no apparent dependence on the pulsar properties [79]. The scaling trend between luminosity and X_{death} , plotted in the right panel of Fig. 7 (a power-law fit to the data has been shown with a red dashed line in order to highlight the trend), is stronger than that of the flux density, but is still minor with respect to the general scatter in the population. As such, one is forced to conclude: (1) that there are many unobserved old pulsars, and (2) the lack of observation does not stem from a limitation of radio telescope sensitivity. Of course neither of these statements are new—the former is natural consequence of pulsar evolution, and the latter has long been appreciated by simulating mock distributions in the pulsar population [65–70]. It is perhaps worth highlighting that the former point could have also been inferred from the mere existence of the MSP population; MSPs are old pulsars [with $\sim \mathcal{O}(\text{Gyr})$ ages [80]] which have been, or are being, spun up by a binary companion—the process of spin-up causes these objects to traverse along fixed trajectories in the P - $\dot{P}$ plane, moving from the death valley to the lower left-hand side of the diagram (and thus the large number of MSPs observed necessitates a large population of unobserved pulsars in the death valley itself).

The role of selection effects in determining biases in the observed pulsar population has been an active area of investigation for many years. In order to understand the potential impact of such effects, we rely on the ability to forward model the full pulsar population—the procedure for simulating mock distributions and the procedure for factoring in observable biases is the subject of the following Appendix.

APPENDIX F: NEUTRON STAR POPULATION SYNTHESIS

As explained in the main text, in order to determine the importance of a pulsar death line in neutron star evolution, one must account for the wide range of selection effects which bias the observed pulsar distribution. The most natural way to account for these is via population synthesis—we outline below the standard procedure for generating mock neutron star populations and highlight the relative importance of each assumption in generating high-quality fits to the observed pulsar distribution.

¹¹The name here is something of a misnomer and should not be interpreted as the height at which radio emission is produced.

1. Spatial distribution

One of the primary ingredients in generating mock pulsar distributions is the determination of their spatial distribution. Neutron stars are born predominantly from core collapse supernova, and thus the distribution of pulsars at birth closely traces that of the stellar distribution in the Galactic disk. As a result of the asymmetry of the core collapse supernova, however, neutron stars are imparted with sizable kick velocities—these initial velocities serve to displace their orbits, leading to an ever-growing discrepancy between the stellar and neutron star distributions. In order to account for this effect, one can sample the location of a neutron star at birth, sample an initial kick velocity, and evolve the neutron star over the course of its lifetime through the Galactic potential. In this work, we adopt a kick distribution

$$P_k(v) = \sqrt{\frac{2}{\pi}} \frac{v^2}{\sigma_k^3} e^{-v^2/\sigma_k^2} \quad (\text{F1})$$

with $\sigma_k = 260$ km/s, which is broadly consistent with the observed proper motions of pulsars [81]. Illustrations of the spatial distributions of a mock neutron star population (produced using the procedure outlined below) are shown in Fig. 8, where the three panels highlight young neutron stars with ages $t \leq 1$ Myr (top), middle-aged neutron stars with ages $t \leq 30$ Myr (middle), and old neutron stars with ages $t \leq 1$ Gyr (bottom). The two panels in Fig. 8 show a top down, and side on, look of the Milky Way. The spiral arm structure is clearly visible in the initial distribution and is washed out over time. The point size and opacity in each panel has been adjusted to ease visibility (the old pulsar population has individual objects extending to large radii, beyond the scale of the Galaxy, and thus the entire region plotted has some lower point density background).

In order to generate the spatial distribution of neutron stars at birth, we follow the procedure which has been adopted in a majority of population synthesis models (see, e.g., Refs. [31,67,82–84]), which involves adopting fits to the observed massive OB stellar distribution (i.e., the stellar distribution leading to the production of neutron stars). In particular, the radial distribution is assumed to follow

$$p(\rho) \propto \left(\frac{\rho}{R_\odot}\right)^{1.9} e^{-5(\frac{\rho}{R_\odot})} \quad (\text{F2})$$

with $R_\odot = 8.5$ kpc, and the height is narrowly confined to the stellar disk, following an exponential distribution

$$p_z \propto e^{-|z|/h_c} \quad (\text{F3})$$

with scale height $h_c = 0.18$ kpc. The angular direction in the $z = 0$ plane is obtained by randomly selecting on the spiral arms, whose loci are described by




FIG. 8. Spatial distribution in the $x - y$ (left) and $\rho - z$ (right) plane of a simulated pulsar population (prior to imposing any selection biases). The various panels highlight the time evolution in the distribution, including only pulsars born within the last 1 Myr (top), 30 Myr (center), and 1 Gyr (bottom). Since there exist strong variations in the number of pulsars satisfying each of these criteria, point size and opacity in each panel has been adjusted in order to make the distributions apparent.

$$\theta_{\ell,i}(\rho) = k_i \log(\rho/\rho_{0,i}) + \theta_{0,i} \quad (\text{F4})$$

with $k_{1/2} = 4.25$ kpc, $k_{3/4} = 4.89$ kpc, $\theta_{0,1} = 1.57$ rad, $\theta_{0,2} = 4.71$ rad, $\theta_{0,3} = 4.09$ rad, $\theta_{0,4} = 0.95$ rad, $\rho_{0,1/2} = 3.48$ kpc, and $\rho_{0,3/4} = 4.90$. Scatter along each arm is included by shifting the pulsar from the central part of the arm by a randomly selected distance r_{raw} , which is drawn

from a normal distribution with mean zero and standard deviation 0.07 kpc.

We evolve pulsars from birth through what is assumed to be an approximately static Galactic potential. The Galactic potential is given by a sum over three components coming from the Galactic nucleus, the bulge, and a combined contribution from the disk and halo,

$$\Phi = \Phi_n + \Phi_b + \Phi_{dh} \quad (\text{F5})$$

with each contribution given by

$$\Phi_b = \frac{-GM_b}{R_b + r}, \quad (\text{F6})$$

$$\Phi_n = \frac{-GM_n}{R_n + r}, \quad (\text{F7})$$

$$\Phi_{dh} = \frac{-GM_{dh}}{\sqrt{(a_G + \sum_i \beta_i \sqrt{z^2 + h_i^2})^2 + b_{dh} + \rho^2}}. \quad (\text{F8})$$

Here, r is the spherical radial coordinate and ρ the cylindrical radial coordinate, and the free coefficients have been fit to the observed stellar motion and are given by: $M_b = 9.3 \times 10^9 M_\odot$, $M_n = 1.0 \times 10^{10} M_\odot$, $M_{dh} = 1.45 \times 10^{11} M_\odot$, $\beta_1 = 0.4$, $\beta_2 = 0.5$, $\beta_3 = 0.1$, $h_1 = 0.325$ kpc, $h_2 = 0.090$ kpc, $h_3 = 0.125$ kpc, $a_G = 2.4$ kpc, $b_{dh} = 5.5$ kpc, $b_b = 0.25$ kpc, and $b_n = 1.5$ kpc [31].

2. Magnetorotational spin-down

As mentioned in the previous sections, the dominant observables for most rotation-powered pulsars are the rotational period and the spin-down rate of the pulsar, with the latter depending on the dipolar magnetic field,¹² the rotational period, and the misalignment angle between the rotational and magnetic axes χ . The distributions that characterize these fundamental properties at birth are not known. As a result, attempts to synthesize mock neutron star populations traditionally adopt reasonable distributions to characterize each of these (typically assuming that their properties are uncorrelated) and fit the free parameters of each distribution using the observed pulsar population. We will follow the same procedure here.

The misalignment angle is often assumed to be randomly selected at birth (i.e., there is no observational or theoretical evidence at the moment to favor a distribution which is preferentially aligned or antialigned), and as such we draw this value from a random distribution,

$$p(\chi_0) = \sin \chi_0 / 2. \quad (\text{F9})$$

¹²Spin-down is driven by electromagnetic energy losses occurring at radii near the light cylinder, $r \sim \Omega^{-1} \gg r_{\text{NS}}$; as such, the spin-down is a probe of the slowest falling multipolar magnetic field component, which is the dipolar component (invalidating this statement would typically require quadrupole field strengths at the surface more than 10^4 times than of the dipolar field strength). We make no attempt in this work to model large higher-order multipolar field components, and note that our results will be broadly consistent so long as such multipoles are not abnormally stronger than the dipolar field.

Magnetic field and period distributions of observed radio pulsars span many orders of magnitude and have no strongly apparent skew, and thus in our baseline analyses we adopt log-normal distributions to characterize both the size of the dipolar field and the rotational period at birth

$$p(B_0) = \frac{1}{\sqrt{2\pi\sigma_{\log B}^2}} e^{-(\log B_0 - \mu_{\log B})^2 / (2\sigma_{\log B}^2)}, \quad (\text{F10})$$

$$p(P_0) = \frac{1}{\sqrt{2\pi\sigma_{\log P}^2}} e^{-(\log P_0 - \mu_{\log P})^2 / (2\sigma_{\log P}^2)}. \quad (\text{F11})$$

These parametrizations are commonly adopted in the literature and have been shown to provide good fits to the observed data [65–70]. Alternative parametrizations of the period distribution have also been proposed (see, e.g., Ref. [85]), and so for the sake of comparison we also adopt a truncated Weibull distribution¹³

$$p(P_0) \propto \left(\frac{P_0}{P_\lambda}\right)^{k-1} e^{-(P_0/P_\lambda)^k} \theta(P_{\text{max}} - P_0) \quad (\text{F12})$$

which serves to show the relative importance of these parametrizations on the inferred dead neutron star population.

The free parameters in Eq. (F10) include the central values and standard deviations in the magnetic field and period ($\mu_{\log B}$, $\sigma_{\log B}$, $\mu_{\log P}$, $\sigma_{\log P}$), and in the case of the Weibull distribution the period is instead characterized by $(k, P_\lambda, P_{\text{max}})$, which capture the skew, the peak, and the truncation scale of the distribution. We fix $P_{\text{max}} = 1$ s, which in practice has a minimal impact on the quality of the fit.

In order to determine the best-fit parameters, one must evolve the population from birth to today. Neutron stars lose their rotational energy (which is dumped into particle acceleration and electromagnetic radiation) at a rate

$$\dot{P} = \beta \frac{B_0^2}{P} (\kappa_0 + \kappa_1 \sin^2 \chi) \quad (\text{F13})$$

where χ is the angle between the rotational period and magnetic field axis (the so-called misalignment angle), the coefficients $\kappa_0 \simeq \kappa_1 \simeq 1$ have been inferred from numerical simulations¹⁴ [86], and $\beta \equiv \pi^2 R_{\text{NS}}^6 / I_{\text{NS}} \simeq 6 \times 10^{-40}$ s/G².

¹³The Weibull distribution has a long tail extending to large periods. This tail has a minimal impact on the fit, but leads to a strong bias in the predicted pulsar birth rate, and thus it is natural in some sense to impose a truncation scale on the initial distribution.

¹⁴Note that standard vacuum dipole radiation can be computed exactly and is proportional to $\sin^2 \chi$. The presence of a corotating plasma that is dragged by the rotating magnetic field, however, yields additional $\mathcal{O}(1)$ corrections to the spin-down rate—while numerical simulations are required to compute these corrections, we note that the energy loss from a plasma-filled force-free magnetosphere is reasonably well constrained.

The process of spin-down drives the neutron star toward a state of alignment, in which $\chi \sim 0$ at a rate given by

$$\dot{\chi} = -\beta\kappa_2 \frac{B^2}{P^2} \sin \chi \cos \chi, \quad (\text{F14})$$

where $\kappa_2 \sim 1$ [86]. Finally, the magnetic field of the neutron star, which is supported by currents inside the star itself, loses energy as those currents slowly dissipate; in general, there are a variety of physical effects which can drive short-term decay of the magnetic field [this includes ambipolar diffusion and the Hall effect, both of which tend to drive a mild suppression of ultrastrong magnetic fields of young pulsars on $\mathcal{O}(\text{kyr})$ timescales] [87–89], but the long-term evolution is believed to be driven by Ohmic dissipation, which is expected to operate on timescales $\tau \gtrsim \mathcal{O}(\text{Myr})$ [89–91]. The efficiency of each of these processes depends on a variety of different factors, including the location of the currents supporting the magnetic field structure, the temperature and composition of the star and crust, the conductivity, etc. Computing the decay of the magnetic field in a self-consistent manner from first principles is an enormously difficult task and remains an active field of research; as a result, we adopt here multiple phenomenological parametrizations of the magnetic field decay, which functional forms and timescales motivated by numerical simulations and pulsar observations [68,92–94], and attempt to demonstrate that any reasonable choice of magnetic field decay that fits the data well will lead to the same observational consequences. Since we are interested predominantly in the late-time evolution of older pulsars, we will neglect the kiloyear evolution of the magnetic field and focus solely on the evolutionary behavior taking place on $t \gtrsim \text{Myr}$ timescales.

One of the most extreme parametrizations one can adopt is the assumption that the magnetic field decays exponentially on timescales τ_{ohm} , i.e.,

$$B(t) = B_0 e^{-t/\tau_{\text{ohm}}}. \quad (\text{F15})$$

This behavior is roughly consistent with the scaling behavior seen in numerical simulations which have crust-confined magnetic fields (an assumption which is unlikely to describe the long-term behavior of the field decay) [68,92]. The value of τ_{ohm} is expected to depend on how the magnetic field penetrates the core, the thickness and temperature of the crust, and the impurity of the crust, with reasonable values of the aforementioned leading to the prediction of Ohmic decay timescales ranging from $\sim \mathcal{O}(\text{Myr})$ to $\sim \mathcal{O}(\text{Gyr})$ (the upper value being the limiting case where magnetic fields do not decay). Observations inferring the kinematic ages of a small subset of the pulsar population [which are far more reliable than the often quoted spin-down timescale, $\tau \sim P/(2\dot{P})$], suggest that long-term exponential decay on timescales $\tau \lesssim 10 \text{ Myr}$ is

not consistent with the observable population [94]. In what follows we will treat τ_{ohm} as a free parameter ranging from 10 Myr to 10 Gyr.

An alternative parametrization has recently been invoked in [84], which takes the form

$$B(t) = B_0 \times \begin{cases} \left(1 + \frac{t}{\tau_1}\right)^{a_1} \left(1 + \frac{t}{\tau_2}\right)^{a_2 - a_1} \times \\ \left(1 + \frac{t}{\tau_{\text{late}}}\right)^{a_{\text{late}} - a_2} \tau_1 < \tau_2 < \tau_{\text{late}} \\ \left(1 + \frac{t}{\tau_1}\right)^{a_1} \times \\ \left(1 + \frac{t}{\tau_{\text{late}}}\right)^{a_{\text{late}} - a_1} \tau_1 < \tau_{\text{late}} < \tau_2 \\ \left(1 + \frac{t}{\tau_{\text{late}}}\right)^{a_{\text{late}}} \tau_{\text{late}} < \tau_1 < \tau_2 \end{cases} \quad (\text{F16})$$

with $\tau_1 = A_1 B_0^{b_1}$, $\tau_2 = A_2 B_0^{b_2}$, $\tau_{\text{late}} = 2 \text{ Myr}$, $A_1 = 10^{14} \text{ yr G}^{-b_1}$, $b_1 = -0.8$, $A_2 = 6 \times 10^8 \text{ yr G}^{-b_2}$, $b_2 = -0.2$, $a_1 = -0.13$, $a_2 = -3.0$, and $a_{\text{late}} = -3.0$. Here, the parameters have been adjusted to fit the decay observed in simulations of magnetic field decay [92] (which again make implicit assumptions about the confinement of magnetic fields to the crust).

At late times, both parametrizations predict that the magnetic field will continue to decay to very small values. As mentioned above, this is unlikely to be a realistic description of the late-time evolution. In order to investigate a potential impact of the unknown late-time evolution, we perform two different types of analyses, one in which these fields are allowed to decay indefinitely and one in which the magnetic field decays until settling into a constant asymptotic value—we choose a final state value by sampling from a log-normal distribution with mean $\mu_{\log_{10} B, \text{final}} = 8.5$ and standard deviation $\sigma_{\log_{10} B, \text{final}} = 0.5$, which is broadly consistent with the properties inferred from millisecond pulsars. These choices are not expected to effect the observable radio pulsar distribution (and thus are expected to fit the data equally well), but have the potential to alter the distribution of dead pulsars (potentially making them relevant when axion hair is present).

As the rotational frequency and magnetic field of the pulsar falls, so does the maximum potential drop across the polar cap $(\Delta V)_{\text{max}} \propto \Omega^2 B$. Consequently, at some point in the evolutionary history of the neutron star, primary particles pulled from the surface of the neutron star cannot be accelerated to sufficiently large boost factors to ignite pair cascades. Since pair cascades are required for the generation of radio emission, the point at which this occurs defines when the pulsar becomes radio silent, or equivalently, when a pulsar “dies.” Furthermore, without pair cascades the magnetosphere becomes “charge starved,” altering the structure of the magnetosphere and the spin-down of the pulsar itself (see, e.g., Ref. [37])—without a dense plasma surrounding the pulsar, the spin-down coefficients tend toward their vacuum

dipole values of $\kappa_0 = 0, \kappa_1 = \kappa_2 = 2/3$. In the following, we will perform two different sets of analyses. In one set of analyses, we will assume pulsar death is abrupt, occurring when pulsars cross our fiducial death line (described in the sections above), leading to an immediate shift in the spin-down coefficients κ_i and making these pulsars unobservable. In the second analysis, we do not impose any notion of a death line, and we allow pulsars to continue evolving with the κ_i values inferred from the force-free limit. The truth likely lies somewhere in between these two analyses, making these well-defined limiting cases.

3. Selection biases and statistical sits

The procedure outlined in the proceeding section tells one how to sample and forward model each pulsar in the mock population, but in order to compare with the observed population one must fold in all observational biases which determine which small fraction of the total population we observed.

We begin by constructing a sample set of $N_{\text{psr,tot}}$ pulsars, where $N_{\text{psr,tot}} = \mathcal{P}(\frac{2\text{psrs}}{\text{cntry}} \tau_{\text{max}})$, with $\mathcal{P}(x)$ being a Poisson distribution with mean x , and τ_{max} chosen to be a sufficiently large number such that all observable pulsars are expected to have an age $\tau \ll \tau_{\text{max}}$ (here we take $\tau_{\text{max}} = 0.2$ Gyr). Ages of each pulsar in the sample are drawn from a flat distribution, and then each pulsar is evolved to determine its properties and location today. We then run through each neutron star in our sample and ask the following questions:

- (i) Is the pulsar dead? If so, and if the presence of a death line is assumed in the analysis, we remove the pulsar from the population. For a default analyses, we adopt the death line shown in Fig. 1. For analyses in which a death line is not imposed, pulsars are retrained with 50% probability—this selection cut is imposed in order to conservatively avoid the question of how bright the radio emission generated on the return current is.
- (ii) Does at least one of the pulsar's radio beams point toward Earth? Here, we randomly select a viewing angle on the sky and use the misalignment angle today χ , alongside the angular opening of the pulsar beam as computed in Eq. (E2), to determine whether part of that beam intersects the angle between the pulsar and Earth. If it does not, the pulsar is removed from the sample.
- (iii) Is the flux density above the detection threshold associated with large-scale pulsar surveys? The signal-to-noise ratio of a typical radio survey can be computed using

$$\frac{S}{N} = \frac{S_{\text{mean}} G \sqrt{n_{\text{pol}} t_{\text{obs}} \delta f_{\text{bw}}}}{\beta T_{\text{tot}}} \sqrt{\frac{P - w_{\text{obs}}}{w_{\text{obs}}}}, \quad (\text{F17})$$

where G is the receiver gain, $S_{\text{mean}} \simeq S_{f,\text{obs}} w_{\text{obs}} / P$ is the mean flux density averaged over a rotational period (with $S_{f,\text{obs}}$ being the observed flux density at Earth), t_{obs} is the observing time, δf_{bw} is the observational bandwidth, $T_{\text{tot}} = T_{\text{rcv}} + T_{\text{sky}}(\theta, \phi)$ is the sum of the receiver temperature and the spatially dependent sky temperature, and w_{obs} is the observed pulse width, which is related to the intrinsic pulse width w_{int} by

$$w_{\text{obs}} = \sqrt{w_{\text{int}}^2 + \tau_{\text{samp}}^2 + \tau_{\text{DM}}^2 + \tau_{\text{scat}}^2}. \quad (\text{F18})$$

Here, the observed width receives additional corrections from the scattering of the radio off the intergalactic medium

$$\tau_{\text{scat}} = 3.6 \times 10^{-9} \text{ DM}^{2.2} (1 + 1.94 \times 10^{-3} \text{ DM}^2), \quad (\text{F19})$$

the sampling time of the telescope τ_{samp} , and pulse smearing

$$\tau_{\text{DM}} = \frac{e^2}{\pi m_e} \frac{\delta f_{\text{ch}}}{f^3} \text{ DM}, \quad (\text{F20})$$

where δf_{ch} is the frequency channel width, and DM is the dispersion measure (i.e., the column density of free electrons along the line of sight). The dispersion measure of each pulsar in the sample can be computed using the Yao-Manchester-Wang Galactic electron model [95]. In order to determine the intrinsic luminosity of the pulsar, we note that the distribution of inferred luminosities from the observed pulsar population roughly follows

$$L = L_0 \dot{P}^\alpha \left(\frac{1 \text{ s}}{P} \right)^{3\alpha} \quad (\text{F21})$$

where $\alpha \sim 0.48$ is a fitting parameter [31,68,84]. In practice, there is scatter about this mean distribution, and thus we sample from a log-normal distribution with mean zero and standard deviation $\sigma_L = 0.8$. The flux at Earth can be directly inferred knowing the distance and the opening angle of each radio beam.

Here, we adopt characteristic values for the radio telescope parameters consistent with what was performed in the broad sky pulsar searches of the Parkes Multibeam Pulsar Survey, Swinburne Intermediate-latitude Pulsar Survey, and High Time Resolution Universe surveys—this includes a signal-to-noise threshold of $S/N = 9$, $\delta f_{\text{ch}} = 3$ MHz, $\tau_{\text{samp}} = 64 \mu\text{s}$, $T_{\text{sky}}(\nu \simeq 1.4 \text{ GHz}) \ll T_{\text{rcv}} = 21 \text{ K}$, $t_{\text{obs}} = 4300 \text{ s}$, $\delta f_{\text{bw}} = 288 \text{ MHz}$, $n_{\text{pol}} = 2$, and $G = 0.735 \text{ K/Jy}$ (see, e.g., compilation of observations in Table 1 of [84]).




FIG. 9. Left: $P - \dot{P}$ diagram showing the comparison between the ATNF catalog (purple) and the best-fit simulated model obtained using an estimated MeerKat flux density threshold and assuming a death line does not exist. In comparison with Fig. 19, which only includes older radio surveys, one can see that the number of pulsars beyond the death line has sizably increased. Right: subsample of distribution of pulsars which are unobservable.

The collection of all pulsars passing these selection cuts defines the observable pulsar sample.

Given a simulated “observable pulsar population,” we are left with the question of how to compare the quality of the mock population with that of the underlying observed population. This is often done by defining a test statistic which characterizes the mismatch between the two populations and then determining the associated p -value by constructing a null distribution, which operates under the hypothesis that in fact both samples are obtained from the same underlying distribution, via bootstrapping. In one dimension, common examples of the test statistic include: the Kolmogorov-Smirnov (KS) test [96], which characterizes the maximum mismatch in the empirically derived cumulative distribution functions (CDFs) of the two datasets (i.e., $\mathcal{D} \equiv \text{Max}[\text{CDF}_X - \text{CDF}_Y]$), the Cramér-von Mises test, which captures the integrated squared difference between CDFs, the binned χ -squared test, which bins the data and looks at the summed differences in the probability density functions (PDFs) of the two distributions, and the energy-distance (ED) metric [97], which computes the expectation values of the pairwise distances between, and within, the various sets (i.e., $\mathcal{D} \equiv 2\mathcal{E}[\|X_i - Y_j\|] - \mathcal{E}[\|X_i - X_j\|] - \mathcal{E}[\|Y_i - Y_j\|]$ where $\mathcal{E}$ is the expectation value, and $\|\cdot\|$ is the Euclidean norm). Each approach has various advantages and disadvantages. For example, the KS and ED tests are intrinsically more sensitive to deviations in the bulk, rather than the tails, of the

distribution (even if the tails are highly incompatible). The χ^2 test, on the other hand, requires the imposition of an arbitrary binning scheme.

In our case, we hope to quantify the extent to which a death line is needed for good fits to the population, and more specifically the extent to which a shifted death line would remain compatible with the observed pulsar distribution. As we will show below, not including a death line in the population models unavoidably leads to the prediction of $\mathcal{O}(100)$ observable pulsars sitting very far below the death line, a number which is largely independent of population modeling assumptions (and a number which grows significantly if one uses a flux density threshold consistent with the currently operating MeerKat telescope, see Fig. 9). This is not an exceptionally large number, contributing to the CDF at only the $\mathcal{O}(2\%)$ level, however, these pulsars are extremely displaced from the bulk of the population and thus would appear as clear outliers would they have been observed. As such, a CDF-based test statistic will not yield sensible results—statistical fluctuations in the bulk distribution of the population yield test statistics which are consistently larger than the highly extended tail (despite the fact that the bulk distributions look extremely similar—see figures below), implying the quality of the fit is in most cases dictated by a combination of unaccounted for noise and small systematics that shift pulsars at central $(P, \dot{P})$, rather than what appears to be a robust incompatible feature of the model.

The solution to this issue is that one must incorporate both the fact that the bulk of the pulsar distribution will have a non-negligible intrinsic scatter across different realizations and the notion that there is some intrinsic cosmic variance (i.e., the observed distribution only represents a single realization). These effects can be naturally included using the χ^2 statistic; here, one can include statistical uncertainties using Poissonian statistics in each bin, as well as correlated features in the PDF itself (which naturally capture the population-to-population scatter) by including off-diagonal elements in the correlation matrix. The size of the off-diagonal elements are not *a priori* known, but since the role of these elements are intended to capture small population-to-population scatter, one can assume that they are distance correlated, and introduce a single tunable parameter λ such that the off-diagonal elements scale as $C_{ij} = \sigma_i \sigma_j e^{-d/\lambda}$ with $d = (i - j)$. As a result, the value of λ is correlated with the adopted binning scheme. In practice, we adopt 35 log-spaced bins across the pulsar parameter space and take a value of $\lambda = 4$, which appears sufficient to wash out small-scale variations without introducing long-range effects (and yields stable results with respect to order-one changes).

Binning and defining PDFs and CDFs in more than one dimension suffers from an intrinsic ambiguity, in that the rescaling of one dimension alters the characteristic value of the distribution function evaluated at the same point. There are often generalized procedures to deal with this issue, see, e.g., Refs. [98,99] for an alternative to the KS test, but one simple procedure to avoid this ambiguity is to project the data onto N_{slic} randomly selected one-dimensional slices, analyze each slice independently, and develop a statistic t_{slic} using, e.g., the maximum of the set, e.g., $t_{\text{slic}} \equiv \text{Max}[\{t_i\}]$. This is the procedure we adopt here, where in practice we take 10^3 random projections in the $(\log_{10} P, \log_{10} \dot{P})$ space. In order to ensure our distance correlator parameter λ is not biased toward variations in $\log_{10} P$ or $\log_{10} \dot{P}$ (which arises since we used a fixed number of bins applied over a different dynamical range in each of the dimensions), we standardize each direction of parameter space using the mean and standard deviation of the distributions inferred from the ATNF catalog. Once the best-fit model is determined, we derive p -values which are normalized to the best-fit model by constructing a null distribution based on bootstrapped χ^2 samples from the dataset and the best-fit model set. This procedure is adopted as the high χ^2 tail of the null distribution can be highly sensitive to the samples used in the bootstrapping procedure, and one can incorrectly infer inflated p -values in scenarios which have significantly worse observed χ^2 statistics.

Note that, when comparing with the ATNF pulsar distribution, we remove pulsars for which $4 \times 10^{10} \text{ G} \leq B \leq B_q$ and $\dot{P} > 0$, i.e., we attempt to remove pulsars which are not spinning down, magnetars, and millisecond pulsars, as these are not intended to be captured by the

population synthesis models. Since we are comparing (at least in some cases¹⁵) with models that impose a sharp pulsar death line, we also remove pulsars in the ATNF catalog which fall below this threshold; in practice, this is only a small number of objects, and we have verified that including these pulsars has effectively no impact on any of the conclusions.

In practice, our analysis works directly on the PDFs and thus do not carry any information on the total number of pulsars in the observable sample—said another way, there is no information in this statistic as to whether the predicted pulsar birth rate is reasonably compatible with theoretical and observational expectations (if, say, ten pulsars were produced in the mock observable pulsar sample, one would have to appeal for birth rates 2 orders of magnitude higher than what is expected theoretically). While the pulsar birth rate is not precisely known, it is expected to follow directly from the theoretically inferred core collapse supernova rate of $\Gamma_{\text{cc}} \sim 1.63 \pm 0.46/\text{century}$ [100]. In general, we perform two sets of analyses, one in which no constraint is imposed on the predicted birth rate and one in which samples are rejected in the inferred birth rate (obtained by computing the population assuming two pulsars are born per century and comparing the size of the predicted mock population with that of the ATNF catalog) exceeds 4/century or falls below 0.5/century. Here, we have chosen sufficiently broad boundaries on the selection thresholds in order to avoid biasing any of our fits.

We provide the full list of models in Table I, highlighting the adopted distributions, evolutionary models, selection cuts, and inferred p -values. Here, we have normalized the p -values to the best-fit model, since an absolute p -value is inherently subject to the unknown level uncertainties, the binning scheme, the treatment of the covariance, etc. (note, however, that in our fiducial analysis, our best-fit models yield p -values ~ 0.20). A number of important trends which can be inferred directly from this table include:

- (1) Models which do not impose a pulsar death line (namely, models 7, 10–17) do not fit the data well, leading to a significant degradation in the inferred p -value. This is not surprising, as all of the models predict $\mathcal{O}(100)$ pulsars sitting well below the death line, a feature which is clearly incompatible in a “by-eye” analysis in Figs. 10–26.
- (2) The initial distribution of rotational periods of pulsars at birth does not strongly favor either the log-normal or the Weibull distribution.
- (3) There is not a strong preference for either magnetic field decay model, however, it is clear that magnetic fields decaying on $\mathcal{O}(\text{Gyr})$ timescales appear to be strongly disfavored.

¹⁵Even though some models do not impose a death line, we remove these pulsars in all population fits in order to ensure consistency between the analyses.

TABLE I. List of parameters used in our neutron star population synthesis. Columns denote model number (an arbitrary label used for identification), the distribution characterizing the dipolar field strength at birth (LN = log -normal), the distribution characterizing the pulsar period at birth (WB = Weibull), the magnetic field decay model [with “Expt” corresponding to Eq. (F15), and “PL” corresponding to Eq. (F16)], whether the late-time evolution of $B(t)$ is truncated at a value determined via the sampling procedure outlined in the text, whether constraints on the pulsar birth rate are applied to the effective likelihood, whether a hard death line is imposed on the simulated population, the best-fit parameters of the initial distribution, and the p -value normalized to the best-fit model p/p_0 .

Model	B_0 distribution	P_0 distribution	$B(t)$	B_{final}	Use Γ_{br}	Death line	Fit parameters	p/p_0
1	LN	LN	Expt. $\tau = 10$ Myr	Yes	Yes	Yes	$(\sigma_{\log B}, \log_{10} B, \sigma_{\log P}, \log_{10} P)$ (0.33, 12.46, 0.73, -0.82)	0.2
2	LN	LN	Expt. $\tau = 10$ Gyr	Yes	Yes	Yes	$(\sigma_{\log B}, \log_{10} B, \sigma_{\log P}, \log_{10} P)$ (0.71, 13.26, 0.14, -1.32)	$< 4 \times 10^{-4}$
3	LN	WB	Expt. $\tau = 10$ Myr	Yes	Yes	Yes	$(\sigma_{\log B}, \log_{10} B, P_{\lambda}, k)$ (0.35, 12.41, 0.39, 58.65)	0.51
4	LN	WB	Expt. $\tau = 10$ Gyr	Yes	Yes	Yes	$(\sigma_{\log B}, \log_{10} B, P_{\lambda}, k)$ (0.41, 12.60, 0.25, 32.82)	$< 4 \times 10^{-4}$
5	LN	LN	Expt. $\tau = 10$ Myr	No	Yes	Yes	$(\sigma_{\log B}, \log_{10} B, \sigma_{\log P}, \log_{10} P)$ (0.40, 12.45, 0.66, -0.71)	0.52
6	LN	LN	Expt. $\tau = 10$ Myr	Yes	No	Yes	$(\sigma_{\log B}, \log_{10} B, \sigma_{\log P}, \log_{10} P)$ (0.40, 12.46, 0.52, -0.98)	0.9
7	LN	LN	Expt. $\tau = 10$ Myr	Yes	Yes	No	$(\sigma_{\log B}, \log_{10} B, \sigma_{\log P}, \log_{10} P)$ (0.37, 12.34, 0.58, -1.01)	$< 4 \times 10^{-4}$
8	LN	LN	PL	Yes	Yes	Yes	$(\sigma_{\log B}, \log_{10} B, \sigma_{\log P}, \log_{10} P)$ (0.36, 12.42, 0.69, -0.82)	1.0
9	LN	WB	PL	Yes	Yes	Yes	$(\sigma_{\log B}, \log_{10} B, P_{\lambda}, k)$ (0.30, 12.40, 0.37, 59.51)	0.19
10	LN	LN	Expt. $\tau = 10$ Myr	No	Yes	No	$(\sigma_{\log B}, \log_{10} B, \sigma_{\log P}, \log_{10} P)$ (0.30, 12.21, 0.40, -0.87)	0.0015
11	LN	WB	Expt. $\tau = 10$ Myr	No	Yes	No	$(\sigma_{\log B}, \log_{10} B, P_{\lambda}, k)$ (0.20, 12.30, 0.49, 44.42)	$< 4 \times 10^{-4}$
12	LN	LN	PL	No	No	No	$(\sigma_{\log B}, \log_{10} B, \sigma_{\log P}, \log_{10} P)$ (0.30, 12.34, 0.70, -1.00)	$< 4 \times 10^{-4}$
13	LN	WB	PL	No	No	No	$(\sigma_{\log B}, \log_{10} B, P_{\lambda}, k)$ (0.61, 12.83, 0.38, 21.97)	$< 4 \times 10^{-4}$
14	LN	LN	Expt. $\tau = 10$ Myr	Yes	No	No	$(\sigma_{\log B}, \log_{10} B, \sigma_{\log P}, \log_{10} P)$ (0.56, 12.70, 0.53, -0.93)	$< 4 \times 10^{-4}$
15	LN	LN	Expt. $\tau = 10$ Gyr	Yes	Yes	No	$(\sigma_{\log B}, \log_{10} B, \sigma_{\log P}, \log_{10} P)$ (0.55, 12.56, 0.30, -0.98)	$< 4 \times 10^{-4}$
16	LN	WB	Expt. $\tau = 10$ Myr	Yes	Yes	No	$(\sigma_{\log B}, \log_{10} B, P_{\lambda}, k)$ (0.45, 12.33, 0.38, 28.42)	$< 4 \times 10^{-4}$
17	LN	LN	PL	Yes	Yes	No	$(\sigma_{\log B}, \log_{10} B, \sigma_{\log P}, \log_{10} P)$ (0.35, 12.47, 0.65, -1.10)	4×10^{-4}




FIG. 10. Comparison of observed and simulated pulsar populations. Left: $P - \dot{P}$ diagram from subselection of the ATNF catalog (purple), compared with the observable pulsar population in model 1. Black line shows the fiducial death line. Right: unobservable part of the mock pulsar population; green points correspond to pulsars which are designated as “active,” but which are not observed due to selection cuts, black points are “dead” pulsars, and red points are active pulsars which would have been designated as “observable” (passing all selection cuts) had one not enforced an abrupt notion of pulsar death. For the sake of simplicity, we only show 5×10^5 samples of the unobservable pulsar population. Note that the gap appearing between the black points and the death line arises because we impose an abrupt shift of the spin-down coefficient κ_1 at pulsar death (which arises from the fact that the spin-down rate of a plasma-filled magnetosphere is larger than that predicted in vacuum)—this assumption is conservative in the sense that it suppresses the number of potentially observable dead pulsars.

- (4) In the case of the best-fitting model, imposing a constraint on the birth rate does not have an appreciable effect (we note that the inferred birth rate tends to larger than one might naively expect and is near the upper threshold of the analysis, although given the uncertainties we hesitate to make any statement beyond this).

The quality of the fit in all cases is highly sensitive to the fit parameters, and thus the posteriors of these parameters are expected to be highly peaked around their best-fit values (although there do exist a few degeneracies in the parameter space which allow for parameters to be shifted in particular directions without spoiling the fit).

Note that we have fit the data using well-known, although slightly outdated, pulsar surveys. The Meerkat telescope has come online in recent years and represents a significant improvement in the state-of-the-art radio instrument, with pulsar physics being one of the fundamental science goals.¹⁶ The Galactic plane survey will lead to significant improvement in the flux density sensitivity and

has already released early science results, including the discovery of new pulsars [102]. We expect this survey to further strengthen the case for a need for a pulsar death line. In order to illustrate this point, we resimulate the population corresponding to model 10, which corresponds to one of the best-fitting models within the class of models which do not impose a death line. Here, we impose a fixed flux sensitivity threshold of $20 \mu\text{Jy}$ in order to simulate the sensitivity of Meerkat. The observed population is shown in Fig. 9 and is found to produce a noticeably worse fit to the observed pulsar population than what is listed in Table I (specifically, we find the p -value degrades by a factor of a few), providing further strength for our conclusions that a death line is a crucial aspect of the pulsar population.

For the sake of completeness, we include figures showing the observed and unobserved pulsar populations for each of the models studied in this work. In the left panels of Figs. 10–26, we compare the observed pulsar population of the ATNF catalog (purple) used in the fits to the mock population of observed pulsars for each model in Table I (gold). The fiducial death line (black) is shown for reference. In the right panel of these figures we instead display a random subset of the unobserved part of the simulated population (the full population is sufficiently

¹⁶The Square Kilometer Array, coming online in the near future, also intends to perform a dedicated pulsar search to improve the understanding of the broader pulsar population [101].




FIG. 11. Same as Fig. 10, but for model 2.




FIG. 12. Same as Fig. 10, but for model 3.

large to be problematic for both storage and figure rendering, and thus this is just done for practical reasons), where the color coding corresponds to: active pulsars (green), dead pulsars that would have been unobservable even if abrupt pulsar death had not been imposed (black), and dead pulsars which are only unobservable because we have assumed they crossed the pulsar death line. Note that,

for models which do not impose pulsar death, all points (by definition) are colored green. Instead, for models imposing pulsar death, the red points highlight potential pulsars that could be revived by the presence of axion hair.

Given that the precise quality of each fit is difficult to judge by eye, we also include in Figs. 27–43 two one-dimensional histograms comparing the predicted best-fit




FIG. 13. Same as Fig. 10, but for model 4.




FIG. 14. Same as Fig. 10, but for model 5.

pulsar population for each model with the observed ATNF population. One- σ statistical error bars are included on the ATNF population for reference. Again, one can notice a

general trend among those populations which do not adopt a death line to overpredict the abundance of low $\dot{P}$ and high P pulsars.




FIG. 15. Same as Fig. 10, but for model 6.




FIG. 16. Same as Fig. 10, but for model 7.




FIG. 17. Same as Fig. 10, but for model 8.




FIG. 18. Same as Fig. 10, but for model 9.




FIG. 19. Same as Fig. 10, but for model 10.




FIG. 20. Same as Fig. 10, but for model 11.




FIG. 21. Same as Fig. 10, but for model 12.




FIG. 22. Same as Fig. 10, but for model 13.




FIG. 23. Same as Fig. 10, but for model 14.




FIG. 24. Same as Fig. 10, but for model 15.




FIG. 25. Same as Fig. 10, but for model 16.




FIG. 26. Same as Fig. 10, but for model 17.




FIG. 27. One-dimensional histograms in P and $\dot{P}$ comparing the observed pulsar population in the ATNF catalog (blue points), as well as the distributions obtained from the best-fit mock population using model 1.




FIG. 28. Same as Fig. 27, but for model 2.




FIG. 29. Same as Fig. 27, but for model 3.




FIG. 30. Same as Fig. 27, but for model 4.




FIG. 31. Same as Fig. 27, but for model 5.




FIG. 32. Same as Fig. 27, but for model 6.




FIG. 33. Same as Fig. 27, but for model 7.




FIG. 34. Same as Fig. 27, but for model 8.




FIG. 35. Same as Fig. 27, but for model 9.




FIG. 36. Same as Fig. 27, but for model 10.




FIG. 37. Same as Fig. 27, but for model 11.




FIG. 38. Same as Fig. 27, but for model 12.




FIG. 39. Same as Fig. 27, but for model 13.




FIG. 40. Same as Fig. 27, but for model 14.




FIG. 41. Same as Fig. 27, but for model 15.




FIG. 42. Same as Fig. 27, but for model 16.




FIG. 43. Same as Fig. 27, but for model 17.

- [1] R. D. Peccei and H. R. Quinn, *CP* conservation in the presence of instantons, *Phys. Rev. Lett.* **38**, 1440 (1977).
- [2] S. Weinberg, A new light boson?, *Phys. Rev. Lett.* **40**, 223 (1978).
- [3] F. Wilczek, Problem of strong *P* and *T* invariance in the presence of instantons, *Phys. Rev. Lett.* **40**, 279 (1978).
- [4] C. Abel *et al.*, Measurement of the permanent electric dipole moment of the neutron, *Phys. Rev. Lett.* **124**, 081803 (2020).
- [5] A. Hook, Solving the hierarchy problem discretely, *Phys. Rev. Lett.* **120**, 261802 (2018).
- [6] L. Di Luzio, B. Gavela, P. Quilez, and A. Ringwald, An even lighter QCD axion, *J. High Energy Phys.* **05** (2021) 184.
- [7] A. Hook and J. Huang, Probing axions with neutron star inspirals and other stellar processes, *J. High Energy Phys.* **06** (2018) 036.
- [8] A. Banerjee, M. A. Buen-Abad, and A. Hook, Stacking the deck: Gambling on a light QCD axion, [arXiv:2507.02049](#).
- [9] R. Balkin, J. Serra, K. Springmann, S. Stelzl, and A. Weiler, White dwarfs as a probe of exceptionally light QCD axions, *Phys. Rev. D* **109**, 095032 (2024).
- [10] K. Bartnick, K. Springmann, S. Stelzl, and A. Weiler, ϕ -dwarfs: White dwarfs probe quadratically coupled scalars, [arXiv:2509.25305](#).
- [11] K. Bartnick, D. Koester, R.-P. Kudritzki, K. Springmann, S. Stelzl, and A. Weiler, New ultralight scalar particles and the mass-radius relation of white dwarfs—The important role of Sirius B, [arXiv:2510.06312](#).
- [12] R. Balkin, J. Serra, K. Springmann, S. Stelzl, and A. Weiler, Heavy neutron stars from light scalars, *J. High Energy Phys.* **02** (2025) 141.
- [13] A. Gómez-Bañón, K. Bartnick, K. Springmann, and J. A. Pons, Constraining light QCD axions with isolated neutron star cooling, *Phys. Rev. Lett.* **133**, 251002 (2024).
- [14] M. Kumamoto, J. Huang, C. Drischler, M. Baryakhtar, and S. Reddy, Pi in the sky: Neutron stars with exceptionally light QCD axions, [arXiv:2410.21590](#).
- [15] Y. Kahn, M. Wentzel, and N. Yunes, Effects of lighter-than-QCD axions on neutron star tidal deformability, [arXiv:2511.13812](#).
- [16] Y. Bai and C. H. de Lima, Electrobaryonic axion: Hair of neutron stars, *J. High Energy Phys.* **05** (2024) 312.
- [17] P. A. Sturrock, A model of pulsars, *Astrophys. J.* **164**, 529 (1971).
- [18] M. A. Ruderman and P. G. Sutherland, Theory of pulsars: Polar gaps, sparks, and coherent microwave radiation., *Astrophys. J.* **196**, 51 (1975).
- [19] K. Chen and M. Ruderman, Pulsar death lines and death valley, *Astrophys. J.* **402**, 264 (1993).
- [20] S. J. Witte, A. Caputo, S. Stelzl, A. Chernoglazov, A. A. Philippov, and S. Rajendran, Axion hair and pulsar electrodynamics: Modeling, discharge dynamics, and particle-in-cell simulations, *Phys. Rev. D* (to be published).
- [21] M. A. Fedderke, J. O. Thompson, R. Cervantes, B. Giaccone, R. Harnik, D. E. Kaplan, S. Posen, and S. Rajendran, Measuring axion gradients with photon interferometry, *Phys. Rev. D* **109**, 015025 (2024).
- [22] M. Pospelov, Odd interaction of axion with matter, *Phys. Rev. D* **58**, 097703 (1998).
- [23] J. E. Moody and F. Wilczek, New macroscopic forces?, *Phys. Rev. D* **30**, 130 (1984).

- [24] C. Gao and A. Stebbins, Structure of stellar remnants with coupling to a light scalar, *J. Cosmol. Astropart. Phys.* **07** (2022) 025.
- [25] R. Balkin, J. Serra, K. Springmann, and A. Weiler, The QCD axion at finite density, *J. High Energy Phys.* **07** (2020) 221.
- [26] P. Di Vecchia and G. Veneziano, Chiral dynamics in the large n limit, *Nucl. Phys.* **B171**, 253 (1980).
- [27] G. Grilli di Cortona, E. Hardy, J. Pardo Vega, and G. Villadoro, The QCD axion, precisely, *J. High Energy Phys.* **01** (2016) 034.
- [28] L. Di Luzio, B. Gavela, P. Quilez, and A. Ringwald, Dark matter from an even lighter QCD axion: Trapped misalignment, *J. Cosmol. Astropart. Phys.* **10** (2021) 001.
- [29] A. Banerjee, J. Eby, and G. Perez, From axion quality and naturalness problems to a high-quality $\mathbb{Z}_{4N}$ QCD relaxation, [arXiv:2210.05690](https://arxiv.org/abs/2210.05690).
- [30] B. Zhang, A. K. Harding, and A. G. Muslimov, Radio pulsar death line revisited: Is PSR J2144–3933 anomalous?, *Astrophys. J.* **531**, L135 (2000).
- [31] C.-A. Faucher-Giguere and V. M. Kaspi, Birth and evolution of isolated radio pulsars, *Astrophys. J.* **643**, 332 (2006).
- [32] S. Konar and U. Deka, Radio pulsar sub-populations (I): The curious case of nulling pulsars, *J. Astrophys. Astron.* **40**, 42 (2019).
- [33] V. Beskin and P. Litvinov, Pulsar death line revisited—I. Almost vacuum gap, *Mon. Not. R. Astron. Soc.* **510**, 2572 (2022).
- [34] V. S. Beskin and A. Y. Istomin, Pulsar death line revisited—II. “The death valley”, *Mon. Not. R. Astron. Soc.* **516**, 5084 (2022).
- [35] A. Philippov, A. Timokhin, and A. Spitkovsky, Origin of pulsar radio emission, *Phys. Rev. Lett.* **124**, 245101 (2020).
- [36] A. Chernoglazov, A. Philippov, and A. Timokhin, Coherence of multidimensional pair production discharges in polar caps of pulsars, *Astrophys. J. Lett.* **974**, L32 (2024).
- [37] A. Philippov and M. Kramer, Pulsar magnetospheres and their radiation, *Annu. Rev. Astron. Astrophys.* **60**, 495 (2022).
- [38] J. Daugherty and A. K. Harding, Electromagnetic cascades in pulsars, *Astrophys. J.* **252**, 337 (1982).
- [39] J. K. Daugherty and A. K. Harding, Pair production in superstrong magnetic fields, *Astrophys. J.* **273**, 761 (1983).
- [40] A. N. Timokhin and A. K. Harding, On the polar cap cascade pair multiplicity of young pulsars, *Astrophys. J.* **810**, 144 (2015).
- [41] A. N. Timokhin and A. K. Harding, On the maximum pair multiplicity of pulsar cascades, *Astrophys. J.* **871**, 12 (2019).
- [42] P. Goldreich and W. H. Julian, Pulsar electrodynamics, *Astrophys. J.* **157**, 869 (1969).
- [43] A. K. Harding, A. G. Muslimov, and B. Zhang, Regimes of pulsar pair formation and particle energetics, *Astrophys. J.* **576**, 366 (2002).
- [44] A. N. Timokhin and J. Arons, Current flow and pair creation at low altitude in rotation powered pulsars’ force-free magnetospheres: Space-charge limited flow, *Mon. Not. R. Astron. Soc.* **429**, 20 (2013).
- [45] A. Caputo, S. J. Witte, A. A. Philippov, and T. Jacobson, Pulsar nulling and vacuum radio emission from axion clouds, *Phys. Rev. Lett.* **133**, 161001 (2024).
- [46] J. D. Jackson, *Classical Electrodynamics* (Wiley, New York, 1998).
- [47] J. Arons, Pair creation above pulsar polar caps: Geometrical structure and energetics of slot gaps, *Astrophys. J.* **266**, 215 (1983).
- [48] M. Kramer and S. Johnston, High-precision geometry of a double-pole pulsar, *Mon. Not. R. Astron. Soc.* **390**, 87 (2008).
- [49] G. Desvignes, M. Kramer, K. Lee, J. van Leeuwen, I. Stairs, A. Jessner, I. Cognard, L. Kasian, A. Lyne, and B. W. Stappers, Radio emission from a pulsar’s magnetic pole revealed by general relativity, *Science* **365**, 1013 (2019).
- [50] S. J. Witte, D. Noordhuis, T. D. Edwards, and C. Weniger, Axion-photon conversion in neutron star magnetospheres: The role of the plasma in the Goldreich-Julian model, *Phys. Rev. D* **104** (2021).
- [51] J. I. McDonald and S. J. Witte, Generalized ray tracing for axions in astrophysical plasmas, *Phys. Rev. D* **108**, 103021 (2023).
- [52] K. Springmann, M. Stadlbauer, S. Stelzl, and A. Weiler, A universal bound on QCD axions from supernovae, [arXiv:2410.19902](https://arxiv.org/abs/2410.19902).
- [53] J. Zhang, Z. Lyu, J. Huang, M. C. Johnson, L. Sagunski, M. Sakellariadou, and H. Yang, First constraints on nuclear coupling of axionlike particles from the binary neutron star gravitational wave event GW170817, *Phys. Rev. Lett.* **127**, 161101 (2021).
- [54] S. J. Witte and A. Mummery, Stepping up superradiance constraints on axions, *Phys. Rev. D* **111**, 083044 (2025).
- [55] A. Caputo, G. Franciolini, and S. J. Witte, Superradiance constraints from GW231123, [arXiv:2507.21788](https://arxiv.org/abs/2507.21788).
- [56] G. L. Smith, C. D. Hoyle, J. H. Gundlach, E. G. Adelberger, B. R. Heckel, and H. E. Swanson, Short range tests of the equivalence principle, *Phys. Rev. D* **61**, 022001 (2000).
- [57] J. Berge, P. Brax, G. Métris, M. Pernot-Borràs, P. Touboul, and J.-P. Uzan, MICROSCOPE mission: First constraints on the violation of the weak equivalence principle by a light scalar dilaton, *Phys. Rev. Lett.* **120**, 141101 (2018).
- [58] D. Noordhuis, A. Prabhu, S. J. Witte, A. Y. Chen, F. Cruz, and C. Weniger, Novel constraints on axions produced in pulsar polar cap cascades, [arXiv:2209.09917](https://arxiv.org/abs/2209.09917).
- [59] J. N. Benabou, C. Dessert, K. C. Patra, T. G. Brink, W. Zheng, A. V. Filippenko, and B. R. Safdi, Search for axions in magnetic white dwarf polarization at Lick and Keck observatories, [arXiv:2504.12377](https://arxiv.org/abs/2504.12377).
- [60] O. Ning and B. R. Safdi, Leading axion-photon sensitivity with NuSTAR observations of M82 and M87, *Phys. Rev. Lett.* **134**, 171003 (2025).
- [61] J. S. Reynés, J. H. Matthews, C. S. Reynolds, H. R. Russell, R. N. Smith, and M. C. D. Marsh, New constraints on light axion-like particles using Chandra transmission grating spectroscopy of the powerful cluster-hosted quasar H1821 + 643, *Mon. Not. R. Astron. Soc.* **510**, 1264 (2021).

- [62] W. C. Haxton, X. Liu, A. McCutcheon, and A. Ray, A continuous galactic line source of axions: The remarkable case of ^{23}Na , *Phys. Rev. Lett.* **135**, 222701 (2025).
- [63] O. Ning, A. Ray, and B. R. Safdi, Axion lines from nuclear de-excitations in galactic stellar populations, *Phys. Rev. D* **113**, 035010 (2026).
- [64] D. F. G. Fiorillo, Á. Gil Muyor, H.-T. Janka, G. G. Raffelt, and E. Vitagliano, Axion-photon conversion in transient compact stars: Systematics, constraints, and opportunities, *J. Cosmol. Astropart. Phys.* **03** (2026) 053.
- [65] A. Lyne, R. Manchester, and J. Taylor, The galactic population of pulsars, *Mon. Not. R. Astron. Soc.* **213**, 613 (1985).
- [66] C.-A. Faucher-Giguere and V. M. Kaspi, Birth and evolution of isolated radio pulsars, *Astrophys. J.* **643**, 332 (2006).
- [67] S. Bates, D. Lorimer, A. Rane, and J. Swiggum, PsrPopPy: An open-source package for pulsar population simulations, *Mon. Not. R. Astron. Soc.* **439**, 2893 (2014).
- [68] M. Gullón, J. A. Miralles, D. Viganò, and J. A. Pons, Population synthesis of isolated neutron stars with magneto-rotational evolution, *Mon. Not. R. Astron. Soc.* **443**, 1891 (2014).
- [69] M. Ronchi, V. Graber, A. Garcia-Garcia, J. A. Pons, and N. Rea, Analyzing the galactic pulsar distribution with machine learning, *Astrophys. J.* **916**, 100 (2021).
- [70] N. Rea, N. Hurley-Walker, C. Pardo-Araujo, M. Ronchi, V. Graber, F. C. Zelati, D. De Martino, A. Bahramian, S. J. McSweeney, T. J. Galvin *et al.*, Long-period radio pulsars: Population study in the neutron star and white dwarf rotating dipole scenarios, *Astrophys. J.* **961**, 214 (2024).
- [71] A. K. Galishnikova, A. A. Philippov, and V. S. Beskin, Simulations of the radio polarization of a precessing pulsar PSR J1906 + 0746, *Mon. Not. R. Astron. Soc.* **497**, 2831 (2020).
- [72] A. Akmal, V. R. Pandharipande, and D. G. Ravenhall, The equation of state of nucleon matter and neutron star structure, *Phys. Rev. C* **58**, 1804 (1998).
- [73] S. E. Gralla, A. Lupsasca, and A. Philippov, Inclined pulsar magnetospheres in general relativity: Polar caps for the dipole, quadrupole and beyond, *Astrophys. J.* **851**, 137 (2017).
- [74] A. M. Beloborodov, Polar-cap accelerator and radio emission from pulsars, *Astrophys. J.* **683**, L41 (2008).
- [75] I. Malov and E. Nikitina, The geometry of radio pulsar magnetospheres, *Astron. Rep.* **55**, 878 (2011).
- [76] B. Zhang, G. Qiao, W. Lin, and J. Han, Three modes of pulsar inner gap, *Astrophys. J.* **478**, 313 (1997).
- [77] A. K. Harding and A. G. Muslimov, Particle acceleration zones above pulsar polar caps: Electron and positron pair formation fronts, *Astrophys. J.* **508**, 328 (1998).
- [78] A. K. Harding and A. Muslimov, Pulsar polar cap heating and surface thermal x-ray emission. 2. Inverse Compton radiation pair fronts, *Astrophys. J.* **568**, 862 (2002).
- [79] S. Johnston, D. Smith, A. Karastergiou, and M. Kramer, The galactic population and properties of young, highly energetic pulsars, *Mon. Not. R. Astron. Soc.* **497**, 1957 (2020).
- [80] B. Kiziltan and S. E. Thorsett, Millisecond pulsar ages: Implications of binary evolution and a maximum spin limit, *Astrophys. J.* **715**, 335 (2010).
- [81] G. Hobbs, D. R. Lorimer, A. G. Lyne, and M. Kramer, A statistical study of 233 pulsar proper motions, *Mon. Not. R. Astron. Soc.* **360**, 974 (2005).
- [82] D. R. Lorimer, A. Faulkner, A. Lyne, R. N. Manchester, M. Kramer, M. McLaughlin, G. Hobbs, A. Possenti, I. Stairs, F. Camilo *et al.*, The Parkes multibeam pulsar survey—VI. Discovery and timing of 142 pulsars and a galactic population analysis, *Mon. Not. R. Astron. Soc.* **372**, 777 (2006).
- [83] M. Cieřlar, T. Bulik, and S. Osłowski, Markov chain Monte Carlo population synthesis of single radio pulsars in the galaxy, *Mon. Not. R. Astron. Soc.* **492**, 4043 (2020).
- [84] V. Graber, M. Ronchi, C. Pardo-Araujo, and N. Rea, Isolated pulsar population synthesis with simulation-based inference, *Astrophys. J.* **968**, 16 (2024).
- [85] S.-S. Du, X.-J. Liu, Z.-C. Chen, Z.-Q. You, X.-J. Zhu, and Z.-H. Zhu, On the initial spin period distribution of neutron stars, *Astrophys. J.* **968**, 105 (2024).
- [86] A. Philippov, A. Tchekhovskoy, and J. G. Li, Time evolution of pulsar obliquity angle from 3D simulations of magnetospheres, *Mon. Not. R. Astron. Soc.* **441**, 1879 (2014).
- [87] A. Passamonti, T. Akgün, J. A. Pons, and J. A. Miralles, The relevance of ambipolar diffusion for neutron star evolution, *Mon. Not. R. Astron. Soc.* **465**, 3416 (2017).
- [88] J. A. Pons and U. Geppert, Magnetic field dissipation in neutron star crusts: From magnetars to isolated neutron stars, *Astron. Astrophys.* **470**, 303 (2007).
- [89] J. A. Pons and D. Viganò, Magnetic, thermal and rotational evolution of isolated neutron stars, *Living Rev. Comput. Astrophys.* **5**, 3 (2019).
- [90] A. Cumming, P. Arras, and E. G. Zweibel, Magnetic field evolution in neutron star crusts due to the Hall effect and ohmic decay, *Astrophys. J.* **609**, 999 (2004).
- [91] D. N. Aguilera, J. A. Pons, and J. A. Miralles, 2D cooling of magnetized neutron stars, *Astron. Astrophys.* **486**, 255 (2008).
- [92] D. Viganò, A. García-García, J. A. Pons, C. Dehman, and V. Graber, Magneto-thermal evolution of neutron stars with coupled ohmic, Hall and ambipolar effects via accurate finite-volume simulations, *Comput. Phys. Commun.* **265**, 108001 (2021).
- [93] S. Mereghetti, J. Pons, and A. Melatos, Magnetars: Properties, origin and evolution, *Space Sci. Rev.* **191**, 315 (2015).
- [94] A. P. Igoshev, Ages of radio pulsar: Long-term magnetic field evolution, *Mon. Not. R. Astron. Soc.* **482**, 3415 (2019).
- [95] J. Yao, R. Manchester, and N. Wang, A new electron-density model for estimation of pulsar and FRB distances, *Astrophys. J.* **835**, 29 (2017).
- [96] V. W. Berger and Y. Zhou, Kolmogorov–Smirnov test: Overview, in *Wiley Statsref: Statistics Reference Online* (John Wiley & Sons, New York, 2014), <https://onlinelibrary.wiley.com/doi/10.1002/9781118445112.stat06558>.
- [97] M. L. Rizzo and G. J. Székely, Energy distance, *Wiley Interdiscip. Rev.* **8**, 27 (2016).
- [98] J. A. Peacock, Two-dimensional goodness-of-fit testing in astronomy, *Mon. Not. R. Astron. Soc.* **202**, 615 (1983).
- [99] G. Fasano and A. Franceschini, A multidimensional version of the Kolmogorov–Smirnov test, *Mon. Not. R. Astron. Soc.* **225**, 155 (1987).

- [100] K. Rozwadowska, F. Vissani, and E. Cappellaro, On the rate of core collapse supernovae in the Milky Way, *New Astron.* **83**, 101498 (2021).
- [101] L. Levin *et al.* (SKAO Pulsar Science Working Group), Understanding the neutron star population with the SKAO telescopes, [arXiv:2512.16156](#).
- [102] P. Padmanabh, E. Barr, S. Sridhar, M. Rugel, A. Damas-Segovia, A. Jacob, V. Balakrishnan, M. Berezhina, M. Bernadich, A. Brunthaler *et al.*, The MPIfR–MeerKAT galactic plane survey–I. System set-up and early results, *Mon. Not. R. Astron. Soc.* **524**, 1291 (2023).