The Astrophysical Journal Letters, 1000:L48 (17pp), 2026 April 1
https://doi.org/10.3847/2041-8213/ae5057

The Radius of PSR J0437–4715 from NICER Data

M. C. Miller[1], A. J. Dittmann[2], I. M. Holt[1], F. K. Lamb[3,4], C. Chirenti[5,6,7,8], Z. Arzoumanian[9], J. Berteaud[5,10], S. Bogdanov[11], K. C. Gendreau[9], W. C. G. Ho[12], S. M. Morsink[13], P. S. Ray[14], R. A. Remillard[15], Z. Wadiasingh[5,7,16], and M. T. Wolff[14]

[1] Department of Astronomy and Joint Space-Science Institute, University of Maryland, College Park, MD 20742-2421, USA; mcmiller@umd.edu
[2] Institute for Advanced Study, 1 Einstein Drive, Princeton, NJ 08540, USA[17]
[3] Illinois Center for Advanced Studies of the Universe and Department of Physics, University of Illinois at Urbana-Champaign, 1110 West Green Street, Urbana, IL 61801-3080, USA
[4] Department of Astronomy, University of Illinois at Urbana-Champaign, 1002 West Green Street, Urbana, IL 61801-3074, USA
[5] Department of Astronomy, University of Maryland, College Park, MD 20742-2421, USA
[6] Astroparticle Physics Laboratory NASA/GSFC, Greenbelt, MD 20771, USA
[7] Center for Research and Exploration in Space Science and Technology, NASA/GSFC, Greenbelt, MD 20771, USA
[8] Center for Mathematics, Computation, and Cognition, UFABC, Santo André, SP 09210-170, Brazil
[9] X-Ray Astrophysics Laboratory, NASA Goddard Space Flight Center, Code 662, Greenbelt, MD 20771, USA
[10] NASA Goddard Space Flight Center, Code 662, Greenbelt, MD 20771, USA
[11] Columbia Astrophysics Laboratory, Columbia University, 550 West 120th Street, New York, NY 10027, USA
[12] Department of Physics and Astronomy, Haverford College, 370 Lancaster Avenue, Haverford, PA 19041, USA
[13] Department of Physics, University of Alberta, Edmonton, AB, T6G 2E1, Canada
[14] Space Science Division, U.S. Naval Research Laboratory, Washington, DC 20375, USA
[15] MIT Kavli Institute for Astrophysics & Space Research, MIT, 70 Vassar Street, Cambridge, MA 02139, USA
[16] Astrophysics Science Division, NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA

Received 2025 December 9; revised 2026 February 26; accepted 2026 March 9; published 2026 March 27

Abstract

Neutron Star Interior Composition Explorer (NICER) data have been used to estimate the masses and radii of the rotation-powered millisecond pulsars PSR J0030+0451, PSR J0740+6620, PSR J0437−4715, PSR J1231−1411, and PSR J0614−3329, sometimes in joint analyses with XMM-Newton data. These measurements provide invaluable information about the properties of cold, catalyzed matter beyond nuclear saturation density. Here, we present the results of our modeling of NICER data on PSR J0437−4715 using several different models of hot thermal X-ray-emitting spots on the stellar surface. For this pulsar, previous Nuclear Spectroscopic Telescope Array observations established that there is also a modulated nonthermal component to the emission, but the previously published analysis of NICER data did not model this component. We find that the Bayesian evidence is significantly higher when the modulated nonthermal component is included and that omission of this component leads to poorer fits to the bolometric NICER data and thus risks bias in the resulting radius estimates. Our models, which we pursue to inferential convergence, therefore have modulated nonthermal emission, and our headline model has in addition three uniform-temperature thermally emitting circular spots. Using this model, the symmetric 68% credible range in the radius is 11.9–15.5 km, which at the independently measured mass of $M = 1.418 \pm 0.044\, M_{\odot}$ is consistent with previous reports of the radius of the $\sim 1.4\, M_{\odot}$ pulsar PSR J0030+0451. We discuss the implications of this measurement for the equation of state of dense matter.

Unified Astronomy Thesaurus concepts: Neutron stars (1108); Neutron star cores (1107); Nuclear physics (2077); X-ray astronomy (1810)

1. Introduction

The equilibrium state of cold matter beyond nuclear saturation density $\rho_{\mathrm{sat}} \approx 2.6 \times 10^{14}\,\mathrm{g\,cm^{-3}}$ cannot be determined from first principles, because of the fermion sign problem (e.g., E. Y. Loh et al. 1990; see Z.-X. Li & H. Yao 2019 for a recent review). Even innovative frameworks such as chiral effective field theory (S. Weinberg 1990, 1991, 1992; C. Drischler et al. 2021) have uncertainties that become large at densities greater than ~ 1.5–$2\,\rho_{\mathrm{sat}}$. Thus, experiments and observations are essential to understanding cold dense matter; and because terrestrial laboratories cannot reach the combination of the high density, temperature much less than the Fermi temperature, and large excess of neutrons over protons seen in neutron stars, observations of neutron stars hold the key to understanding nuclear physics in this extreme state.

Over the last decade or so, several lines of astronomical observations have contributed constraints on the macroscopic properties of neutron stars and, thus, on the microscopic equation of state (EOS; for cold, catalyzed neutron stars, this is conveniently represented as the pressure P as a function of energy density ϵ). These lines include: the discovery of high-mass neutron stars by radio timing (P. B. Demorest et al. 2010; J. Antoniadis et al. 2013; H. T. Cromartie et al. 2020; E. Fonseca et al. 2021; A. Saffer et al. 2025), including the more model-dependent masses inferred from the so-called "black widow" pulsars (e.g., R. W. Romani et al. 2022, 2026) and the binary neutron star coalescence GW170817, which provided an

---

[17] NASA Einstein Fellow.

upper limit to the tidal deformability of $\sim 1.4\,M_\odot$ neutron stars (B. P. Abbott et al. 2017; S. De et al. 2018).

An important contribution has been made by X-ray measurements of the neutron stars PSR J0030+0451 (M. C. Miller et al. 2019; T. E. Riley et al. 2019; S. Vinciguerra et al. 2024), PSR J0740+6620 (M. C. Miller et al. 2021; T. E. Riley et al. 2021; A. J. Dittmann et al. 2024; T. Salmi et al. 2024a), PSR J0437−4715 (D. Choudhury et al. 2024a), PSR J1231−1411 (T. Salmi et al. 2024b), and PSR J0614−3329 (L. Mauviard et al. 2025), using Neutron Star Interior Composition Explorer (NICER) and XMM-Newton X-ray data complemented by radio measurements.

To this list, we can add potentially more model-dependent inferences from the electromagnetic afterglow of GW170817, the general properties of short gamma-ray bursts (A. Bauswein et al. 2013; C. L. Fryer et al. 2015; S. Lawrence et al. 2015; B. Margalit & B. D. Metzger 2017; E. R. Most et al. 2018; L. Rezzolla et al. 2018; M. Ruiz et al. 2018; C. Chirenti et al. 2019, 2023; V. Guedes et al. 2025), and other X-ray data (M. C. Miller & F. K. Lamb 2016; F. Özel et al. 2016; A. L. Watts et al. 2016; J. Nättilä et al. 2017), however many of these methods are dominated by systematic errors and their results therefore cannot be used with confidence (M.C. Miller & F. K. Lamb 2016). By 2030, it may be possible to measure the moment of inertia of PSR J0737−3039A via pulsar timing at radio wavelengths to a precision of $\sim 10\%$ (H. Hu et al. 2020).

Here, we report our analysis of NICER data on the 173.7 Hz pulsar PSR J0437−5715 (see W. Becker & J. Trümper 1993, 1999 for the discovery and initial characterization of the X-ray pulsations). Like the other pulsars observed with NICER, PSR J0437−4715 is not actively accreting. Instead, the X-ray pulsations evident in NICER data are caused by hot spots on the surface of the spinning neutron star, formed when highly relativistic (characteristic Lorentz factors likely $\gtrsim 10^7$) beams of electrons and positrons, which are generated as part of the process that produces coherent radio emission, penetrate the surface and deposit their energy (A. K. Harding & A. G. Muslimov 2011). PSR J0437−4715 is in a binary system, so radio observations place tight constraints on the mass ($M = 1.418 \pm 0.044\,M_\odot$), the distance ($D = 156.96 \pm 0.11$ pc, from the I. S. Shklovskii (1970) effect on the orbital period derivative), and the inclination angle of the observer to the orbital axis ($i = 137.506 \pm 0.016°$—all numbers from D. J. Reardon et al. 2024).

Because PSR J0437−4715 has the highest X-ray flux among nonaccreting millisecond pulsars (with, e.g., ~ 3 times the NICER count rate of PSR J0030+0451 and ~ 30 times the NICER count rate of PSR J0740+6620), models of the X-ray hot spots are put to a particularly exacting test for this pulsar. We therefore explore three models of the thermally emitting hot spots (all of which use nonmagnetic model atmospheres of fully ionized hydrogen) using the NICER data (see Section 3). Our models also include contributions that do not vary in a way commensurate with the pulsar rotational frequency and thus are unmodulated in the folded pulse waveform, based on a model of background counts from R. A. Remillard et al. (2022) and from an angularly nearby active galactic nucleus (AGN). We further include a modulated nonthermal contribution, which we model as a power law, based on the observed modulation in higher-energy Nuclear Spectroscopic Telescope Array (NuSTAR) data (S. Guillot et al. 2016).

The observation of pulsed emission with NuSTAR in the energy range between 2 and 20 keV (S. Guillot et al. 2016) is significant, because it is implausible that received flux with this spectrum and energy range could be produced by thermal emission originating on the surface of the neutron star. Furthermore, S. Guillot et al. (2016) showed that a power-law emission model describes both the NuSTAR data and the lower-energy (0.5–2.0 keV) XMM-Newton observations (S. Bogdanov 2013) of a nonthermal spectral component. This implies that a small fraction of the pulsed emission appearing in the NICER data (0.3–3.0 keV) originates from a region corotating with the neutron star, such as the magnetosphere. We model this emission with a modulated power-law component, as described in Section 3.6. As shown in Section 4, our models that include this modulated power law result in much better Bayesian evidence than the models that attempt to describe all of the NICER emission as arising from thermally emitting hot spots on the star's surface.

Our spot models are two uniform-temperature circular spots, two uniform-temperature oval spots, and three uniform-temperature circular spots. The spots have independent angular sizes and temperatures, can be anywhere on the star, and can overlap or not, as the fit prefers. As we discuss in Section 5, when the modulated power law is not included, the bolometric fit to the data is poorer than it is when the modulated power law is included. When the modulated power law is included, all three of the models with a modulated power law have similar Bayesian log evidence, which is several tens larger than when the modulated power law is not included. Thus, the modulated power law must be included, but none of the spot models are preferred on the basis of their Bayesian evidence. The three-circle model has the highest log likelihood; its radius posterior is the broadest of the three models (although the radius posterior overlaps at the 90% credible level for all models). We therefore feature this model. Its symmetric 68% credible range for the radius is 11.9–15.5 km.

We have also compared our best NICER-only model with the XMM-Newton data on PSR J0437−4715. We find that extra XMM-Newton background flux (up to double the predicted counts at low energies, with a smaller fractional excess at higher energies), beyond that measured in the immediate vicinity of the pulsar, is necessary to obtain agreement between our model and the XMM-Newton data. The source of this extra emission is not clear; it could be a source along the projected line of sight, or an intrabinary shock, or something else. We therefore focus entirely on fits of the NICER data, rather than attempting joint fits with the XMM-Newton data.

In Section 2, we describe our selection of data and the calibration of the instruments. In Section 3, we give a brief description of our analysis methods; more details may be found in previous papers (M. C. Miller et al. 2019, 2021; A. J. Dittmann et al. 2024). In Section 4, we discuss our models and compare our best fit with the data, finding no systematics in the residuals. In Section 5, we present our results, and in Section 6, we discuss the implications that our PSR J0437−4715 radius measurements, combined with other results, have for the EOS of cold, catalyzed matter above $\rho_{\rm sat}$. In Section 7, we present our conclusions, and in several appendices we provide details of our fits including corner plots. Posterior samples from our analysis are available in Zenodo via doi:10.5281/zenodo.17833896.

2. Observations

For this analysis, we make use of NICER X-ray Timing Instrument (XTI) exposures of PSR J0437–4715, covering the interval between 2017 July 6 and 2021 July 31. The NICER data of this pulsar contain a significant flux contribution from the bright Seyfert II AGN RX J0437.4−4711 (J. P. Halpern & H. L. Marshall 1996), situated 4$'\!.$18 away. To minimize contamination due to the AGN and other neighboring sources, NICER observations of PSR J0437–4715 were carried out using a pointing position 1$'\!.$47 to the southwest of the pulsar (see S. Bogdanov et al. 2019a for further details).

To produce a clean event list suitable for analysis, we adopt the 3C50 model approach, primarily because it is designed to also produce a background spectrum from the source dataset with reliable uncertainty estimates in the energy band of interest. The details of the 3C50 model are described in R. A. Remillard et al. (2022).

Valid exposure segments for PSR J0437−4715 were first determined with the HEASoft tool `nimaketime`, and the times were masked by ±30 s at the orbital day/night boundary. From there, continuous intervals longer than 200 s were selected for further cleaning. For each good time interval (GTI), the raw NICER spectrum is extracted, a background prediction is obtained with the 3C50 model, and six filtering tests (see the next paragraph) are performed on the results. Only GTIs with 50 detectors active are included in the final selection, with events from detectors with `DET_ID` 14 and 34 always excluded, as they frequently exhibit elevated count rates well above the average of the other detectors.

The 3C50 background model makes use of three nonsource count rates from the on-source GTIs of the actual observations being considered to select and rescale background components from a reference library of models. Two filters apply to parameters needed for the background model: the slow chain noise rate normalized to 50 Focal Plane Modules at 0.0–0.25 keV, $nz < 220$ counts s^{-1}, and the normalized rate of good events at 15–18 keV (where the performance of XTI is such that effectively no astrophysical source signal is expected), $ibg < 0.2$ counts s^{-1} (nz, ibg, and hbg are defined in R. A. Remillard et al. 2022). The next two filters act on the net background-subtracted spectrum, $S0_{\rm net} < 0.15$ counts s^{-1} and $hbg_{\rm net} < 0.05$ counts s^{-1}, where the corresponding energy bands are 0.2–0.3 keV for $S0$ and 13–15 keV for hbg. Finally, when the net pulsar spectrum shows very faint or undetectable intensity above 2 keV, filters are added in the C band (2–4 keV), $C_{\rm net} < 0.1$ counts s^{-1}, and D band (4–12 keV), $D_{\rm net} < 0.3$ counts s^{-1}. This rigorous filtering procedure reduces the initial unfiltered 2.736 Ms exposure to 1.310 Ms.

NICER also conducted on-axis observations of the AGN RX J0437.4−4711 during 2017 July 10–20, 2021 October 8–10, December 10–26, and 2022 February 28–March 19, in order to establish its flux and range of variability. The mean spectrum of the AGN was determined for each epoch using the 3C50 background model and Level 2 filtering was applied to each GTI. The average intensity (0.3–3.0 keV) was determined to be 6.26 counts s^{-1}, while the rms deviation between the four epochs is 25%. The expected contribution from the AGN in the PSR J0437–4715 dataset is represented by the mean AGN spectrum, scaled down by a factor of 42, corresponding to the reduction in the telescope effective area based on the angular separation between the AGN and the offset pointing direction.

The final selection of filtered NICER events was phase-folded, based on the ephemeris from a publicly available Parkes Pulsar Timing Array radio timing solution (see D. J. Reardon et al. 2024) using the PINT[18] (J. Luo et al. 2021) `photonphase` tool.

3. Methods

Our approach to the generation of waveforms and the analysis of NICER and XMM-Newton data has been described in several previous papers. We summarize our methods, with references to papers with more details, in the following subsections.

3.1. Spacetime

To generate waveforms, we select heated patterns on a neutron star surface and trace the emitted photons through vacuum to the observer using the "oblate Schwarzschild approximation" (S. M. Morsink et al. 2007; M. AlGendy & S. M. Morsink 2014; for an update, see Z. Jakab & S. M. Morsink 2025). In this approximation, we model the surface as an oblate spheroid with a shape given by the rotation rate and the mass and radius the star would have without rotation, and we include all special relativistic effects. However, we use the spherical Schwarzschild spacetime for the external spacetime of the star. For rotation rates below ∼600 Hz, the error introduced by this approximation, compared with the numerically exact solution obtained by solving the full Einstein field equations for a rigidly rotating star, is several times smaller than the statistical uncertainties in the data (S. Bogdanov et al. 2019b). We have also verified the accuracy of our waveform codes against other codes (S. Bogdanov et al. 2019b, 2021; D. Choudhury et al. 2024b).

3.2. Magnetic Fields

We assume that the magnetic field in the neutron star atmosphere is too weak to affect the atmospheric structure or the spectrum or beaming of photons from the surface, at NICER energies. In practice, this is a good approximation when the electron cyclotron energy is much less than the 0.3 keV minimum energies that we analyze—i.e., when the surface magnetic field is much less than $\sim 3 \times 10^{10}$ G (for more details, see S. Bogdanov et al. 2021). PSR J0437−4715 has a rotation period of $P = 0.005757$ s and an observed period derivative $\dot{P} = 5.729 \times 10^{-20}$ (J. P. W. Verbiest et al. 2008; D. J. Reardon et al. 2024). However, the majority of this observed period derivative comes from proper motion (the I. S. Shklovskii 1970 effect), and when the measured proper-motion values etc. from D. J. Reardon et al. (2024) are used to calculate the magnitude of this effect, we find that the intrinsic period derivative is only $\dot{P}_{\rm intrinsic} = 1.369 \times 10^{-20}$. Inserting this into Equation (12) of I. Contopoulos & A. Spitkovsky (2006) implies that for a magnetic field inclination θ of $\pi/2$ and a field configuration with $\alpha = 0$ (i.e., so that the closed region of the magnetic field reaches the light cylinder), the surface field strength is $B \approx 4 \times 10^8$ G, which is weak enough to be neglected unless significantly stronger multipoles exist on the surface. See Section 3.1 of M. C. Miller et al. (2021) and Section 3.1.3 of A. J. Dittmann et al. (2024) for more details about this argument.

---

[18] https://github.com/nanograv/PINT

3.3. Atmosphere Models

Because the neutron star surface gravity is so strong and the star is not actively accreting, we expect that the atmosphere will consist purely of the lightest element present (C. Alcock & A. Illarionov 1980). Given the 5.74 days orbital period of the PSR J0437−4715 binary (J. P. W. Verbiest et al. 2008), which implies an orbital radius that is easily large enough to have contained a donor with a hydrogen envelope, the atmosphere is likely to be pure hydrogen. We attempted one run using a helium atmosphere and found that the resulting fits were poor ($\chi^2/\mathrm{dof} = 10103.4/8350$ for three circles with no modulated power law; by comparison, as we show in Section 5.1, the best fit of the same model with a hydrogen atmosphere has $\chi^2/\mathrm{dof} = 8371.1/8350$). Therefore, our results focus on fully ionized nonmagnetic hydrogen atmospheres, using the NSX code (W. C. G. Ho & D. Lai 2001) for the model atmospheres.[19]

We work in the deep-heating approximation: that the energy emerging from the photosphere was effectively deposited infinitely deep in the atmosphere by the high-energy electrons and positrons that hit the surface as part of the pulsar process. This is a good approximation at the Lorentz factors $\gamma \sim 10^7$ expected for millisecond pulsars (A. K. Harding & A. G. Muslimov 2002). If, contrary to expectations, most of the deposited energy comes from far lower-energy electrons and positrons, $\gamma < 100$, then the energy is deposited in shallower layers. In that case, M. Bauböck et al. (2019) and T. Salmi et al. (2020) found, as expected, that the resulting beaming pattern of emission would be broader than in the deep-heating approximation. This would lead to a lower modulation fraction in the waveform. To match an observed modulation fraction, other parameters would therefore need to adjust. In particular, a lower compactness $GM/(Rc^2)$ (for a star of gravitational mass M and circumferential radius R) leads to a higher modulation fraction with all else equal, so if the energy is deposited in a shallow layer but the data are analyzed assuming deep heating, one would expect that the inferred radius would be smaller than the actual radius.

T. Zhao et al. (2025) found the opposite (that the inferred radii would be too large if the beaming pattern of emission is broader than assumed in analyses), but given that in their synthetic data they used (1) 60 times as many spot counts as in the PSR J0740+6620 data, with no background, and (2) a restricted spot distribution (two identical antipodal spots; NICER pulsars do not have antipodal spots), it is not clear that this result is applicable to the stars studied using NICER.

3.4. Surface Emission Patterns

Our hot spots on the stellar surface are constructed with either uniform-temperature circular spots or, in a generalization of that model, uniform-temperature oval spots. The spots can overlap in an arbitrary way; we number our spots and assume that a pixel in more than one spot emits with the effective temperature of the lowest-numbered spot that includes the pixel. We do not expect that real heated regions on neutron stars come in these simple patterns, but previous work (F. K. Lamb et al. 2009a, 2009b; K. H. Lo et al. 2013; M. C. Miller & F. K. Lamb 2015; I. M. Holt et al. 2025) suggests that even when models deviate from the true pattern, if there is a statistically good fit, then the inferred radius is not significantly biased. This is understood to be because the angularly broad emission from a thermally emitting point on the surface combined with gravitational light deflection blur details of the spot patterns.

3.5. Nonpulsar NICER Counts

In addition to the NICER counts produced by photons from the spots, there will in general be counts produced in a number of other ways, including from the particle background, optical loading from the Sun, other X-ray sources in the field (whether resolved or unresolved), or unassociated X-ray emission from the pulsar system (such as from a pulsar wind nebula or from intrabinary shocks). What these other counts have in common is that none of them are expected to be modulated at a frequency commensurate with the rotational frequency of the pulsar. Our most general and conservative treatment of these “background” counts in NICER data therefore assumes that in each NICER pulse-invariant (PI) channel, independently, there is an additional component that is independent of rotational phase, and in our Bayesian analysis we marginalize over that component (see Section 3.4 of M. C. Miller et al. 2019 for more details). In some cases, we may have reliable information about one or more components of the background. In the case of PSR J0437−4715, the 3C50 data selection procedure discussed in Section 2 also outputs a background estimate, which we implement as discussed in Section B.1 of T. Salmi et al. (2022). There is also an AGN in the field of view, which we incorporate in our background estimates.

3.6. The New Component: Modulated Nonthermal Pulsar Emission

An additional element that must be added in our analysis of PSR J0437−4715 is the modulated nonthermal component, discovered in NuSTAR data by S. Guillot et al. (2016). Because this component is modulated at the pulsar rotational frequency, it cannot be incorporated into the unmodulated background. For simplicity, D. Choudhury et al. (2024a) did not include this component, but we show that when it is not included, the bolometric fit is poorer than when the modulated nonthermal component is included, and the inferred radius could be biased. We have found that a good representation of the modulation is a Gaussian profile in the rotational phase. We therefore model the photon number flux from the power law, at an energy E and rotational phase ψ, using the expression

$$F_{\rm PL}(E, \psi) = N_{\rm PL}(E/1\,{\rm keV})^{-\alpha_{\rm PL}}[1+ A_{\rm PL}\exp(-(\psi - \psi_{\rm PL})^2/(2(\Delta\psi_{\rm PL})^2))]. \quad (1)$$

This model describes the modulated nonthermal component using five parameters: the overall normalization $N_{\rm PL}$, the power-law index $\alpha_{\rm PL}$, the fractional modulation $A_{\rm PL}$, the rotational phase $\psi_{\rm PL}$ of the maximum modulation, and the width $\Delta\psi_{\rm PL}$ of the Gaussian. The priors for the power-law parameters are listed in Table 1; the ranges for $\alpha_{\rm PL}$ and $N_{\rm PL}$ were motivated by the results in Table 2 of S. Guillot et al. (2016). However, we used a much larger range in the

---

[19] Partial ionization is possible for the lower-temperature spots, but previous work has found little difference in the inferred radius between fully and partially ionized hydrogen atmospheres (see, e.g., section 3.1 of M. C. Miller et al. 2021), and the tables for partially ionized hydrogen are not as complete as for fully ionized hydrogen.

Table 1
Primary Parameters of the Pulse Waveform Models Considered in This Work

Parameter	Definition	Assumed Prior
$c^2R_e/(GM)$	Inverse of stellar compactness	3.2–8.0
$M/M_\odot$	Gravitational mass	$\mathcal{N}(1.418, 0.044)$
θ_{c1} (rad)	Colatitude of spot 1 center	0 to π
$\Delta\theta_1$ (rad)	Spot 1 half-extension	0–3
$kT_{\mathrm{eff},1}$ (keV)	Spot 1 effective temperature	0.011–0.5
f_1	Spot 1 elongation factor	$\log_{10} f_1$ flat from 0 to +1
ψ_1 (rad)	Spot 1 tilt angle	0–π
$\Delta\phi_2$ (cycles)	Spot 2 longitude difference	0–1
θ_{c2} (rad)	Colatitude of Spot 2 center	0 to π
$\Delta\theta_2$ (rad)	Spot 2 half-extension	0–3
$kT_{\mathrm{eff},2}$ (keV)	Spot 2 effective temperature	0.011–0.5
f_2	Spot 2 elongation factor	$\log_{10} f_2$ flat from 0 to +1
ψ_2 (rad)	Spot 2 tilt angle	0–π
$\Delta\phi_3$ (cycles)	Spot 3 longitude difference	0–1
θ_{c3} (rad)	Colatitude of Spot 3 center	0 to π
$\Delta\theta_3$ (rad)	Spot 3 half-extension	0–3
$kT_{\mathrm{eff},3}$ (keV)	Spot 3 effective temperature	0.011–0.5
f_3	Spot 3 elongation factor	$\log_{10} f_3$ flat from 0 to +1
ψ_3 (rad)	Spot 3 tilt angle	0–π
$f_{\mathrm{back,3C50}}$	Multiplier of 3C50 background spectrum	$\mathcal{N}(1, 0.0461)$
$f_{\mathrm{back,AGN}}$	Multiplier of AGN background spectrum	0.5–1.5
N_{PL} (keV^{-1} cm^{-2} s^{-1})	Power-law normalization	$\log_{10} N_{\mathrm{PL}}$ flat from −7.0 to −2.0
α_{PL}	Power-law photon index	$\mathcal{N}(1.5, 0.25)$
A_{PL}	Power-law modulation fraction	0–1
ψ_{PL} (cycles)	Power-law phase	0–1
$\Delta\psi_{\mathrm{PL}}$ (cycles)	Power-law width	0.01–0.49
θ_{obs} (rad)	Observer inclination	$\mathcal{N}(0.74166, 0.000279)$
N_H (cm^{-2})	Neutral H column density	(0–20) × 10^{20}
d (kpc)	Distance	$\mathcal{N}(0.15696, 0.00011)$
f_{eff}	NICER effective area factor	$\mathcal{N}(1, 0.1)$

Note. These are the parameters and priors for our most complex model (Model 6). This model has three uniform-temperature oval spots with arbitrary overlap, plus a modulated power law with a Gaussian phase profile, to take into account the modulated nonthermal component visible in the NuSTAR data on PSR J0437−4715, plus prior information on the NICER background from the model of R. A. Remillard et al. (2022) and from the nearby AGN. Simpler models (e.g., with two oval spots or two or three circular spots) have the same priors for the parameters they have in common with the three-oval model. Except where noted, the prior is flat over the given range. $\mathcal{N}(a, b)$ means a Gaussian with mean a and standard deviation b.

normalization parameter than the range considered by S. Guillot et al. (2016), because the NuSTAR data are at far larger energies than the NICER data, and we therefore wanted to guard against possible errors in energy extrapolation.

3.7. Statistical Methods

Our overall statistical inference approach is Bayesian: we explicitly specify the models and their priors, use Bayes' Theorem to obtain the posteriors, and marginalize over nuisance parameters (such as background parameters) to obtain the posteriors for parameters of interest, such as the mass and radius. Because the parameter space is complex, with 20–30 parameters in the runs we discuss in Section 5, it is necessary to use sophisticated statistical samplers to obtain posteriors. In the next subsection, we discuss our sampling procedure in detail. See Section 4 of M. C. Miller et al. (2019) for tests of our sampling with synthetic two-circle and two-oval data similar to the NICER data on PSR J0030+0451, which were performed using the `MultiNest` sampler (F. Feroz et al. 2009) rather than the `pocomc` sampler (M. Karamanis et al. 2022) that we feature here.

3.8. Bayesian Methodology

We take a two-stage approach to evaluating models for the NICER data on PSR J0437−4715, similar to our analyses of the J0030+0451 and J0740+6620 data (M. C. Miller et al. 2019, 2021; A. J. Dittmann et al. 2024): in the first stage, we sample each model in a relatively agnostic manner, exploring the entire prior volume and estimating the Bayesian evidence for each model; in the second stage, we confirm or correct the results of the first stage using Markov Chain Monte Carlo (MCMC) methods, which are less suited to efficient global exploration but have more reliable convergence properties. In the first stage, we employed the `pocomc` sampler (M. Karamanis et al. 2022), which combines sequential Monte Carlo sampling (see, e.g., C. A. Naesseth et al. 2019 for a review) with normalizing flows (e.g., G. Papamakarios et al. 2021), to achieve relatively robust and efficient sampling and evidence estimation without prior knowledge about which subdomains of the prior volume ought to be sampled. We follow up the initial `pocomc` analyses using the `emcee` sampler (D. Foreman-Mackey et al. 2013) to confirm or correct the inferred posterior distributions before we use them in EOS inference. We find that the posterior after resampling with `emcee` is essentially indistinguishable from the posterior after a converged `pocomc` run.

3.8.1. Sequential Monte Carlo Analyses

We begin our analyses using the publicly available `pocomc` package,[20] which utilizes normalizing flows to precondition a given target distribution, before sampling it using an adaptive sequential Monte Carlo scheme (M. Karamanis et al. 2022). For our purposes, the crucial benefits of this scheme are that: (1) it is able to estimate the Bayesian evidence for model comparison; (2) it samples from the entire prior volume without relying on potentially biasing initialization schemes; and (3) when applied to our NICER inferences, it tends to outperform many packages that implement variations on nested sampling, which offer the aforementioned benefits (J. Skilling 2004). Specifically, the nested sampling algorithm MultiNest (F. Feroz et al. 2009) has been used extensively in previous NICER analyses (e.g., M. C. Miller et al. 2019, 2021; T. E. Riley et al. 2019, 2021; D. Choudhury et al. 2024a; A. J. Dittmann et al. 2024; T. Salmi et al. 2024a). Although MultiNest has a tendency to produce biased posterior and evidence estimates (see J. Buchner 2016, 2023; B. E. Nelson et al. 2020; P. Lemos et al. 2023; A. Dittmann 2024), it has provided useful starting points for more robust MCMC analyses (M. C. Miller et al. 2019, 2021; A. J. Dittmann et al. 2024). However, the additional model complexity necessary to describe both the thermal emission from the surface of PSR J0437–4715 and the modulated power-law emission, presumably from its magnetosphere, proved too expensive for the methods used previously.[21]

Sequential Monte Carlo sampling (e.g., L. Stewart & P. McCarty 1992; N. Gordon et al. 1993; G. Kitagawa 1996; C. A. Naesseth et al. 2019) lies at the heart of `pocomc`, which is itself a generalization of importance sampling. Importance sampling by itself uses a known importance sampling density ($\rho(\theta)$, where θ represents the parameters) to guide the estimation of an unknown target density $p(\theta)$, which can expedite inference if the importance sampling density is similar to the target density but can otherwise degrade performance severely: one cannot generally construct a useful importance sampling density a priori. Sequential Monte Carlo constructs a sequence of sampling and target densities, in this application sampling first from the prior $\pi(\theta)$ to estimate an annealed posterior $p_i(\theta) \propto \pi(\theta)\mathcal{L}(\theta)^{\beta_i}$, where $\mathcal{L}(\theta)$ is the likelihood, for a sequence of β_i ranging from 0 and 1, sampling from the prior to the posterior. By gradually adjusting β_i, the importance sampling density at each stage can be smoothly adapted so that the next stage functions properly. For this additional effort, one gains an estimate of the Bayesian evidence and the ability to sample efficiently from the posterior once the corresponding importance sampling density has been determined, while avoiding any need to make assumptions at the outset about which subsets of the prior to sample.

MCMC methods can struggle with correlated or skewed target distributions. This can be ameliorated by preconditioning—a process by which a complicated target distribution is transformed to a distribution that is simpler and more amenable to sampling. `pocomc` uses normalizing flows (e.g., G. Papamakarios et al. 2021) to accomplish this task, which are generative models that, over the course of sampling, infer bijective mappings between a given distribution and a Gaussian distribution. Preconditioning accelerates the sequential Monte Carlo algorithm by roughly an order of magnitude when sampling nontrivial target distributions (M. Karamanis et al. 2022).

While modeling PSR J0437−4715, we found that `pocomc` provides significantly more robust results than, for example, MultiNest analyses of comparable computational cost. Still, given the complicated high-dimensional likelihood surfaces encountered in our analyses, our `pocomc` analyses demanded significant convergence testing. The thoroughness of a given `pocomc` analysis is controlled by $N_{\rm effective}$—the effective number of weighted particles used to explore the parameter space at each stage of the algorithm. While the default value is set to $N_{\rm effective} = 2^9 = 512$, we found that values of $N_{\rm effective} \gtrsim 2^{17}$ were necessary to thoroughly sample our most complicated models, in the sense that doubling the resolution no longer significantly affected the inferred log evidence values and the widths of various posterior distributions.

We consider a `pocomc` analysis to have converged when two analyses of the same dataset and model, differing by a factor of at least 2 in $N_{\rm effective}$, produce estimates of the log evidence and log maximum likelihood that are consistent within the uncertainties estimated within `pocomc` (typically <0.01). Furthermore, when considering nested models with differing complexities (e.g., two spots versus three spots), we only consider the more complicated analysis to be converged, judging how thoroughly it has explored the parameter space, if it results in a maximum likelihood *at least* as high as that found by the simpler model.

We emphasize, however, that compared to other (often more expensive) sampling methods, `pocomc` can occasionally underestimate posterior widths (e.g., J. U. Lange 2023; M. J. Williams et al. 2025), although not to the same extent as MultiNest (e.g., M. C. Miller et al. 2019, 2021; J. Ih & E. M. R. Kempton 2021; A. J. Dittmann et al. 2024; T. Salmi et al. 2024a; I. M. Holt et al. 2025; see M. Hoogkamer et al. 2025 for a recent comparison between MultiNest and UltraNest). Thus, for our models that include the necessary power-law components, we have followed up each `pocomc` analysis with additional MCMC sampling to ensure the accuracy of our reported posteriors. We now discuss this procedure in more detail.

3.8.2. MCMC Analyses

Following each of our `pocomc` analyses, we used the estimated posterior distributions to draw initial walker positions for an MCMC analysis using the `emcee` package (D. Foreman-Mackey et al. 2013), employing the affine-invariant “stretch” proposal of J. Goodman & J. Weare (2010).

---

[20] https://github.com/minaskar/pocomc

[21] Notably, MultiNest is based on rejection sampling, so the cost of the analyses grows exponentially with the number of model parameters. Algorithms based on Monte Carlo techniques have better (polynomial) asymptotic scaling. Nested sampling packages that can employ slice sampling, which has better scaling, include PolyChord (W. J. Handley et al. 2015), dynesty (J. S. Speagle 2020), and UltraNest (J. Buchner 2021), but we have found `pocomc` (M. Karamanis et al. 2022) better suited to our analyses.

Because we utilized a proposal distribution that satisfied detailed balance, the distribution of walkers will eventually converge to sample from the target distribution directly. In practice, because the results produced by `pocomc` were very nearly converged, we typically observed that the tails of each posterior broadened slightly, by a few percent, after which we continued drawing samples to reduce discretization errors in our posterior estimates.

Each analysis used $2^{13} = 8192$ walkers, for which we drew initial positions by resampling a Gaussian kernel density estimate of the corresponding `pocomc` posterior distributions. We judged the convergence of each analysis by monitoring the 1st, 16th, 50th, 84th, and 99th percentiles of each distribution and identifying the point at which each stopped exhibiting secular variations in time. We terminated each analysis after accruing $\sim 10^6$ (effectively) independent samples for each model. See Appendix B for a comparison of the `pocomc` radius posterior with the final radius posterior after sampling with `emcee`, for our featured model with three circular spots and a modulated power law.

4. Models Used in Our Analysis

We employ several models in our analysis of the NICER data on PSR J0437−4715. Table 1 has a description of the primary parameters for our most complex model (in terms of number of parameters), which has three uniform-temperature oval spots, as well as a modulated power law to represent the modulated nonthermal component. Less complicated models (with two circular spots, two oval spots, or three circular spots) have the same priors as the three-oval model for the parameters they have in common.

In our fits, we incorporate information about known components of the unmodulated NICER background (see T. Salmi et al. 2022 for a similar use of NICER background information in the analysis of PSR J0740+6620). In the case of PSR J0437−4715, this includes not just the “3C50" NICER instrumental background model of R. A. Remillard et al. (2022) but also counts from the angularly nearby AGN RX J0437.4−4711 (J. P. Halpern & H. L. Marshall 1996). For both sources of information, we produced a smoothed count spectrum using a cubic spline. We then added to the phase-channel model produced from the spots a background that was an energy-independent factor times the nominal background count spectrum; the priors on the factors are listed in Table 1. Because of our assumption that the background counts are not modulated at a frequency commensurate with the rotation frequency of PSR J0437−4715, we distribute these extra counts uniformly in rotational phase.

In addition to these known sources of backgrounds, we allow for unknown additional sources of background in the NICER data. As described in more detail in Section 3.4 of M. C. Miller et al. (2019), we do this independently in each NICER PI channel and marginalize over the added background by fitting a Gaussian to the likelihood near the peak and performing an analytic integration of the marginalization integral. Thus, in most of our fits, there is a lower limit to the total background counts in each NICER PI channel based on known backgrounds, but because we use a Gaussian prior on the 3C50 contribution this is not a hard limit.

In more detail, the models we fit are as follows. In each case, the emission from all spots uses a fully ionized nonmagnetic hydrogen atmosphere, and we fit only to the NICER data rather than performing a joint NICER−XMM fit. Moreover, in all models but one, the spots can be anywhere on the star, with or without overlaps. All of our models include background contributions from the 3C50 model of R. A. Remillard et al. (2022) as well as the angularly nearby AGN. Our six models are summarized in Table 2 and described in more detail below:

Table 2

Description of the Pulse Waveform Models Considered in This Work

Model Number	Spot Model	Modulated PL?	Converged?
1	Two circles	No	Yes
2	Three circles	No	Yes
3	Two circles	Yes	Yes
4	Two ovals	Yes	Yes
5	**Three circles**	**Yes**	**Yes**
6	Three ovals	Yes	No

Note. In every model, we incorporate estimates of the NICER background from the empirical 3C50 model (R. A. Remillard et al. 2022) and the angularly nearby AGN, and there is no maximum set on the background. Moreover, all models use fully ionized, deep-heating, nonmagnetic hydrogen for the atmosphere, and in each case we analyze only the NICER data rather than performing a joint NICER−XMM analysis. All models assume uniform-temperature spots. The bold font highlights our featured model.

1. Two circular spots with no modulated power law;
2. Three circular spots with no modulated power law;
3. Two circular spots with a modulated power law;
4. Two oval spots with a modulated power law;
5. Three circular spots with a modulated power law—Model 5 is our featured model; and
6. Three oval spots with a modulated power law—this model did not converge.

Models 1–5 appeared to converge, based on a comparison of independent `pocomc` runs with different values of the precision parameters. However, at the next level of complexity (Model 6, with three oval spots with a modulated power law), we were not able to establish convergence. For example, at the highest resolution that we were able to employ in the three-oval run, the maximum log likelihood was ∼8 less than the maximum log likelihood for the three-circle run, even though three circles is a subset of three ovals. We therefore stopped our model development at Model 5.

5. Results of the Analysis of the NICER Pulse Waveform Data

5.1. Quality of Fits and Inferred Radii

Table 3 summarizes the qualities of the best fit of each of the models that we consider, along with the $\pm 1\sigma$ range of the inferred radius in each model. We see that all of the models have phase-channel χ^2 probabilities high enough for the fit to be acceptable. However, we also see that the maximum log likelihood and log evidence are much worse for the models (1 and 2) that do not have a modulated power-law component. The remaining three models overlap in their $\pm 1\sigma$ radius intervals. Based on the log evidences, we will use only the

Table 3
Summary of the Results of the Analysis Using Our Pulse Waveform Models

Model Number	$\Delta \ln \mathcal{L}_{\max}$	$\Delta \ln \mathcal{Z}$	Phase-channel χ^2/dof (Prob)	Bolometric χ^2/dof (Prob)	$\pm 1\sigma$ $R_{\rm eq}$ (km)
1	−48.93	−25.31	8359.14/8354 (0.482)	36.13/27 (0.112)	12.11–13.76
2	−38.69	−36.56	8371.1/8350 (0.4331)	47.97/27 (0.008)	10.25–12.66
3	−8.76	−0.40	8352.72/8349 (0.49)	38.92/27 (0.0644)	13.60–15.96
4	−8.15	−1.48	8350.76/8345 (0.48)	39.64/27 (0.055)	13.94–16.10
5	0	0	**8335.00/8345 (0.53)**	**33.39/27 (0.18)**	**11.88–15.49**

Notes. The maximum likelihoods and evidences in the second and third columns are relative to our featured model—i.e., our Model number 5, highlighted in bold font. The phase-channel and bolometric χ^2 values are for the maximum-likelihood parameter combination for each model. The effective number of parameters affecting the bolometric light curve is much smaller than the total number of parameters and is mostly independent of the model. We have found empirically, using synthetic data, that the bolometric χ^2 for the maximum-likelihood parameter combination for 32 phases is distributed as if it had ~27 degrees of freedom and thus, effectively, five important parameters, for all of the listed models. The maximization of the phase-channel log likelihood is essentially the same as the minimization of the phase-channel χ^2, but if the goal were to minimize the bolometric χ^2, then the distribution would be different. For example, for parameters that give a log likelihood within ~10 of the maximum, the minimum bolometric χ^2 for Model 1 is 36.13 (i.e., it coincides with the maximum-likelihood parameter combination), the minimum bolometric χ^2 for Model 2 is 43.44, and the minimum bolometric χ^2 for Model 5 (our featured model) is 20.56. Note that even if all five parameters used to describe the bolometric power law were to count against the bolometric degrees of freedom for those fits, then $\chi^2/\mathrm{dof} = 20.56/22$ for the best bolometric fit using Model 5, which would thus still have a higher bolometric probability than the models that do not include a modulated power law.

three-circle posterior when we discuss the implications of our results for the EOS of cold, catalyzed matter beyond nuclear saturation density.

Figure 1 compares the radius posteriors for our converged models with a modulated power law—i.e., Models 3, 4, and 5. Although the posteriors are not identical, they have substantial overlap. This demonstrates that, at least when a modulated power law is included, the posterior is not sensitive to the details of the spot model.

See Appendix C for the full posteriors for our fit with three circular spots and a modulated power law (Model 5).

5.2. Characteristics of the Best-fit Model with Three Circular Spots and a Modulated Power Law

In this section, we explore the characteristics of the best fit we obtained to the NICER PSR J0437−4715 data using our featured model—i.e., Model 5—which has three circular spots plus a modulated power law with a Gaussian phase profile.

Figure 2 shows the spot distribution in the best fit. Two of the spots are close to each other and in the same hemisphere as the observer, albeit with very different temperatures, whereas the third spot, which has a much larger solid angle, is in the other hemisphere. This could be a configuration that approximates one spot with a spatially asymmetric distribution of temperatures and one with a more uniform temperature distribution.

Figure 3 shows the residuals, which we compute using

$$\chi = (\mathrm{data} - \mathrm{model})/(\mathrm{model})^{1/2} \tag{2}$$

for the best fit and the data for each NICER PI channel (vertical axis) and the rotational phase (horizontal axis). No patterns are visible, and no individual phase-channel bin has a significantly larger $|\chi|$ than would be expected with this number of bins. This, in addition to the overall $\chi^2/\mathrm{dof} = 8335/8345$ (a probability of 0.53 if the model is correct), shows that this test does not raise any concerns.

Figure 4 shows the bolometric best-fit model and data (upper panel) and residuals (again using Equation (2)), as a function of rotational phase. As with the phase-channel χ, in addition to the fit being formally acceptable ($\chi^2/\mathrm{dof} = 33.39/27$, which has a probability of 0.18 if the model is correct), there are no evident patterns or strong outliers, and thus this test does not suggest any problems with the fit. The bolometric and phase-channel comparisons between the best fit and the data yield one-way tests: if the fit were poor, we would need to examine the models closely, but the fact that the fits are acceptable does not guarantee that the model is correct.

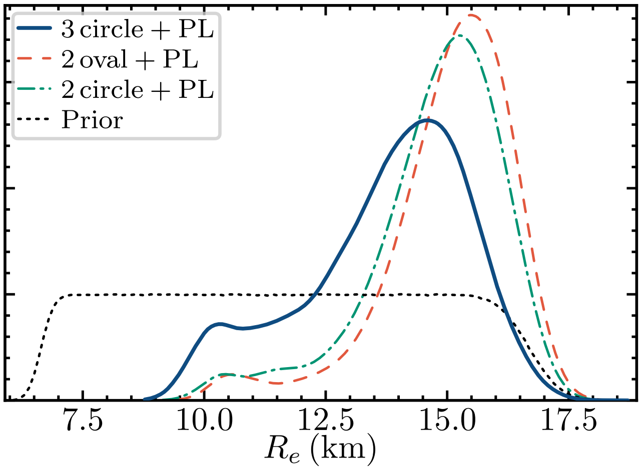

Figure 1. Radius posteriors for each of our three models with a modulated power law (the vertical axis is linear in the probability density); from top to bottom in the legend, the model numbers are 5, 4, and 3. The dotted line shows the prior. There is a large overlap of the posteriors of all three models, which shows that among the models with a modulated power law, the posterior is not especially sensitive to the model chosen.

Figure 5 shows the mass–radius posterior for our fit using three circular spots plus a modulated power law—i.e., Model 5. The dotted curve on the one-dimensional mass plot (top left) shows the prior from the radio observations of D. J. Reardon et al. (2024), from which it is clear that the radio observations provide most of the information about the mass. The bottom right panel shows the one-dimensional radius posterior and demonstrates that it is bimodal with one mode predominating.

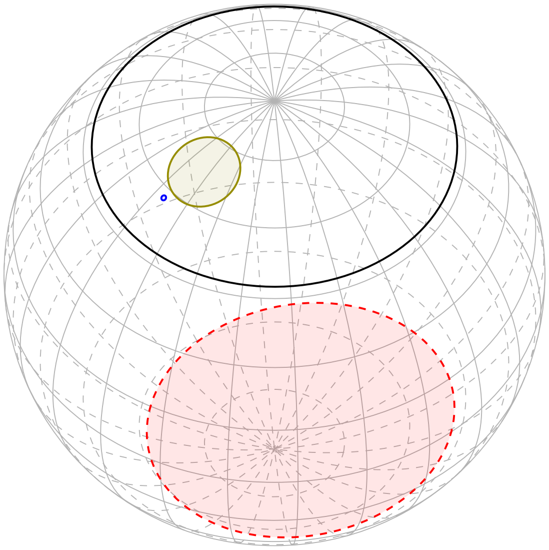

Figure 2. Spot locations, sizes, and temperatures for the best fit of our featured model with three uniform-temperature circular spots plus a modulated power law—i.e., Model 5. The smallest spot has an effective temperature, as measured by a comoving observer on the surface, of 0.17 keV; the middle-sized spot has an effective temperature of 0.031 keV; and the largest spot has an effective temperature of 0.11 keV. The solid black circle indicates the colatitude of the observer, 0.742 radians, which is strongly constrained by radio observations. See Appendix C for the full set of parameter values for this best fit. Bearing in mind that the temperature distribution is surely not actually uniform circles, this could be an indication that the two spots in the hemisphere of the observer represent a single spot with a range of temperatures.

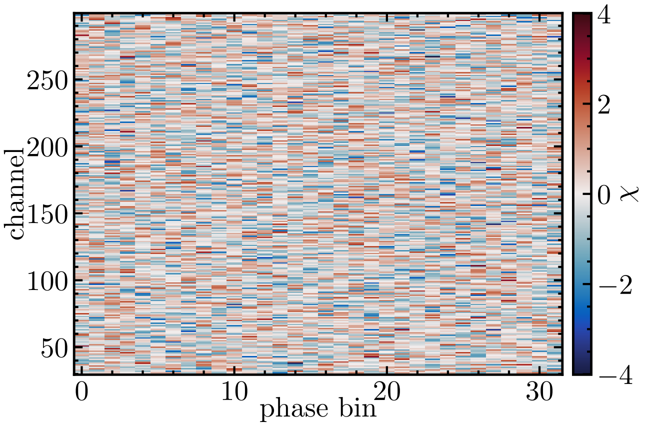

Figure 3. Phase-channel residuals for the best fit to the data of our model with three uniform-temperature circular spots plus a modulated power law, over the full set of 32 uniformly spaced rotational phases and NICER PI channels 30–299 inclusive. Here, for a phase-channel bin i, if the data have d_i counts and the model predicts m_i counts, we define $\chi \equiv (m_i - d_i)/\sqrt{m_i}$. In addition to the fit having an overall acceptable $\chi^2/\text{dof} = 8335.00/8345$ (for which the probability of this χ^2 or higher for a correct model is 53%), there are no patterns visible and no individual phase-channel bins with an unexpectedly large $|\chi|$. This is a one-way test: our satisfactory fit does not guarantee that the model is correct, but a very low probability would indicate that we would need to look more closely at our model.

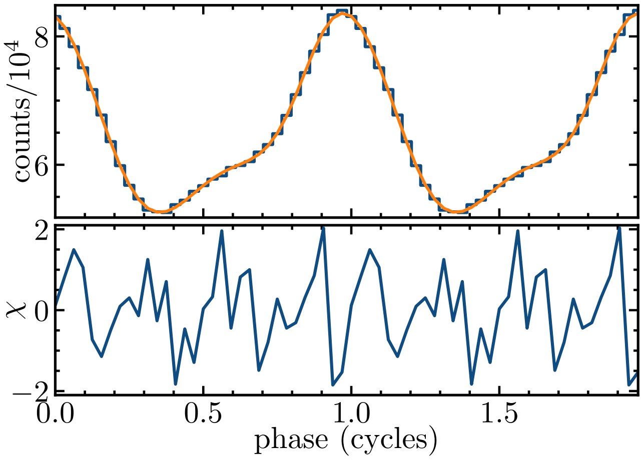

Figure 4. Top: comparison between the bolometric data (histogram) and the best model with three uniform-temperature circular spots plus a modulated power law (red line). Bottom: residuals between the data and the best model. Here, we plot two full rotational cycles, with 32 uniformly distributed rotational phases per cycle. In addition to the overall adequate bolometric χ^2/dof = 33.39/27 (18% probability if the model is correct), we see no strong outliers and no obvious patterns in the residuals. As with the phase-channel χ, this is a one-way test, which might detect a strong deviation between the model and the data but cannot guarantee the model's correctness if it passes the test.

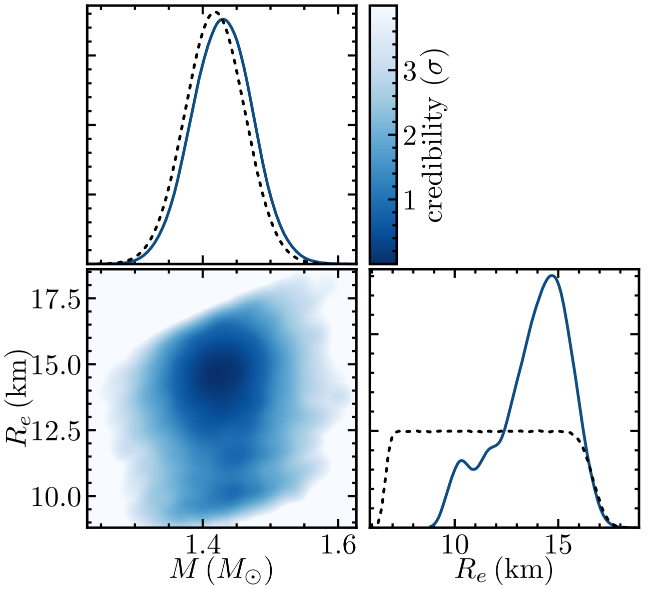

Figure 5. Mass and radius posteriors, using our best fit to the data of our model with three uniform-temperature circular spots plus a modulated power law—i.e., our Model 5. The dotted lines show the priors. The approximate upper diagonal boundary in the mass–radius plot corresponds to our $c^2R_e/(GM) = 8$ prior upper limit, whereas the approximate lower boundary is set by the likelihood rather than by our prior $c^2R_e/(GM) = 3.2$ lower limit. The mass posterior is single-peaked and is shifted only slightly from the tight prior given by radio observations, whereas the radius posterior is bimodal (with a larger mode at higher radii).

5.3. Comparison with D. Choudhury et al. (2024a)

We now compare these results with the earlier analysis of NICER observations of PSR J0437−4715 by D. Choudhury et al. (2024a), which analyzed the same dataset that we have analyzed, using the X-PSI code (T. E. Riley et al. 2023), and did not include a modulated power-law component. More specifically, we performed runs using the single-temperature (ST) + protruding double-temperature (PDT) configuration of D. Choudhury et al. (2024a; first introduced in T. E. Riley et al. 2019). ST+PDT is a restricted subset of a full three-circle configuration, in which (1) none of the three spots are allowed to have angular radii greater than $\pi/2$; (2) Spot 1 is not allowed to overlap with Spot 2 or Spot 3; and (3) Spots 2 and 3 must overlap, at least at a point. For this configuration and with no constraints on the background (labeled "No BKG" in Figure 6 of D. Choudhury et al. 2024a), D. Choudhury et al. (2024a) obtained a $\pm 1\sigma$ radius range of 9.81–10.46 km. This is not the headline model of D. Choudhury et al. (2024a), and the headline model uses more precise `MultiNest` settings (a sampling efficiency of 0.3 and 20,000 live points compared with 0.3 and 2000 for the ST+PDT run), but the handling of the background is closer to one of our options, and thus this provides a convenient point of comparison.

The radius range we find for the no-background ST+PDT configuration with no modulated power law, using parallel-tempered `emcee` (see Section 3.5 of M. C. Miller et al. 2021 for methodological details), is 10.16–12.21 km at $\pm 1\sigma$, which is much broader than reported by D. Choudhury et al. (2024a). In previous parameter estimations for other pulsars, such as PSR J0740+6620 (A. J. Dittmann et al. 2024; T. Salmi et al. 2024a), our radius posteriors have typically been wider than those produced by X-PSI analyses, which is probably due to differences in the statistical sampling of the two methods. Our best fit for our version of the ST+PDT model has an acceptable phase-channel $\chi^2/\mathrm{dof} = 8375.25/8352$ (a probability of 0.427), but the bolometric $\chi^2/\mathrm{dof} = 47.97/27$ has a probability of only 0.008 and therefore might indicate a poor fit.

A comparison of our results, with and without a modulated power law, and the results of D. Choudhury et al. (2024a; which do not include a modulated power law) demonstrates that in addition to yielding a better bolometric fit, the power-law component shifts the posterior to higher radii. Thus, if the modulated power law is not included, it risks bias in the radius posterior.

6. Implications for the EOS

6.1. Statistical Method

As in M. C. Miller et al. (2019, 2021) and A. J. Dittmann et al. (2024), we base our EOS inference on the method described in M. C. Miller et al. (2020). This is a fully Bayesian approach in which we begin with a selection of EOSs (three families of which are described below), which by fiat all have the same prior probability. The (unnormalized) posterior probability of each EOS is then updated by calculating the product of the likelihoods of several datasets: these include the high masses of PSR J1614−2230 ($M = 1.937 \pm 0.014\,M_\odot$; G. Agazie et al. 2023) and PSR J0348+0432 ($M = 1.806 \pm 0.037\,M_\odot$; A. Saffer et al. 2025); the tidal deformability of two double-neutron-star coalescences seen at ground-based gravitational-wave detectors (B. P. Abbott et al. 2017, 2020); the mass and radius measurements of PSR J0030+0451 (M. C. Miller et al. 2019) and PSR J0740+6620 (A. J. Dittmann et al. 2024); and now our mass and radius measurement of PSR J0437−4715. At densities up to half of nuclear saturation (i.e., up to roughly the crust–core transition density; see K. Hebeler et al. 2013), we assume that the QHC19 (G. Baym et al. 2019) EOS applies, but because uncertainties at densities this low do not contribute significantly to uncertainties in the neutron star mass and radius, other choices of the low-density EOS give very similar results.

We perform our likelihood calculations for the maximum mass, radius as a function of mass, and tidal deformability as a function of mass by assuming that the star is rotating much slower than its maximum possible rate (typically ∼1500–2000 Hz for most EOSs at $1.4\,M_\odot$), and thus that the Tolman–Oppenheimer–Volkoff (TOV) equations (J. R. Oppenheimer & G. M. Volkoff 1939; R. C. Tolman 1939) are accurate. They are highly accurate for the 173.7 Hz rotation rate of PSR J0437−4715, as can be seen in Figure 13 of M. C. Miller et al. (2019), which demonstrates that the difference between the equatorial radius of a nonrotating star and a star rotating at 200 Hz is much smaller than the measurement uncertainty, for a wide range of gravitational masses and EOSs.

If we wish to compare a measured mass with the maximum stable gravitational mass of a nonrotating star, then we need only to specify the EOS. In contrast, if we need to compute the likelihood of a NICER mass–radius posterior with the expectations from a given EOS, then we must also assume a probability distribution for the central density of the star, given that this is the boundary condition for the TOV and related equations. Here, we follow the convention in M. C. Miller et al. (2021), in which, for a given EOS, the prior on the central density is quadratic between the central density ρ_{min} of a $1\,M_\odot$ star and the maximum central density ρ_{max} of a stable star. That is, if x is distributed uniformly between 0 and 1, then the central density prior is given by $\rho_c = \rho_{\mathrm{min}} + x^2(\rho_{\mathrm{max}} - \rho_{\mathrm{min}})$, except that we skip densities yielding unstable stars (this can happen for $\rho_c < \rho_{\mathrm{max}}$ in, for example, EOSs with phase transitions, which are included in our sample). This contrasts with the linear central density prior chosen in M. C. Miller et al. (2019), but the two density priors lead in practice to very similar posteriors.

When we compute the likelihood of a NICER mass–radius posterior, or of a gravitational-wave measurement of the tidal deformability Λ as a function of neutron star mass, we need to integrate over the central densities. For each central density, we thus need to determine the probability of the observation given an EOS model. Because we have discrete samples of posterior points in the M–R or M–Λ planes, it is necessary to employ a procedure that uses these discrete samples to estimate a smoothed probability distribution in two dimensions. As in M. C. Miller et al. (2019) and M. C. Miller et al. (2021), we use Gaussian kernel density estimation for this purpose, but with 0.1 times the standard bandwidth recommended by B. W. Silverman (1986). See Section 5.1 of M. C. Miller et al. (2021) for additional details.

We also include a constraint of $S = 32 \pm 2$ MeV (M. B. Tsang et al. 2012) on the nuclear symmetry energy S (defined as the difference in the total energy per nucleon between matter with an equal density of protons and of neutrons, $E_{\mathrm{SNM}} \approx -16$ MeV, and the total energy per nucleon of pure neutron matter) at nuclear saturation density. M. C. Miller et al. (2021) approximated S as $S = \epsilon/n - m_n c^2 - E_{\mathrm{SNM}}$ at number density n, which is equivalent to assuming that the matter at nuclear saturation density is pure neutrons. Here, we compute S at saturation density by imposing beta equilibrium on the matter, including the possible presence of muons, using the formulae in Section II of D. Blaschke et al. (2016). This makes a few percent difference in S but, as noted in M. C. Miller et al. (2021), the inclusion of a constraint on S has only a minor effect on the EOS posteriors, because it applies at comparatively low densities.

Some studies of the EOS of high-density matter apply additional theoretical constraints, based on, e.g., chiral effective field theory (see C. Drischler et al. 2021 for a recent review) or perturbative quantum chromodynamics (see T. Gorda et al. 2023 for a recent example, but also see R. Kumar et al. 2024; D. Mroczek 2024; D. Mroczek et al. 2024). Here, we do not include such considerations, so that our samples can be used with different assumptions about nuclear physics.

6.2. EOS Models

We follow A. J. Dittmann et al. (2024) in using Gaussian processes to extend our EOS beyond half of nuclear saturation density (although we note that other choices are possible, such as piecewise polytropes or spectral forms; see M. C. Miller et al. 2019, 2021). Gaussian processes were introduced in EOS estimates by P. Landry & R. Essick (2019) and developed extensively and mostly by the same authors (e.g., R. Essick et al. 2020a, 2020b, 2023; P. Landry et al. 2020; I. Legred et al. 2021, 2024; E. Finch et al. 2025). See C. E. Rasmussen & C. K. I. Williams (2006) for an overall introduction to Gaussian processes. The basic idea is: starting with a fiducial function $f_0(\boldsymbol{x})$ tabulated at particular values $\boldsymbol{x}$ of the independent variable, one can create variations $f(\mathbf{x}) = f_0 + \Delta f(\boldsymbol{x})$ on that function by drawing $\Delta f(\boldsymbol{x})$ from a multivariate Gaussian distribution. In the context of EOS estimation, following L. Lindblom (2010) and P. Landry & R. Essick (2019), our independent variable is the log pressure, and the dependent variable is

$$\phi \equiv \ln\left(c^2 \frac{d\epsilon}{dP} - 1\right), \tag{3}$$

which is constructed so that (as is necessary for a full Gaussian distribution) physically possible values of ϕ can range from $-\infty$ (corresponding to the boundary of causality, with a sound speed $c_s = (dP/d\epsilon)^{1/2} = c$) to $+\infty$ (corresponding to the boundary of stability $c_s = 0$). For consistency with M. C. Miller et al. (2021), we choose as our fiducial function $\phi_0 = 5.5 - 2.0(\log_{10} P - 32.7)$ (where the pressure P is in cgs units). This matches well the EOSs listed at the CompOSE website https://compose.obspm.fr/ but also drives the sound speed to values approaching the speed of light at several times nuclear saturation density. More flexible choices are possible; see, for example, R. Essick et al. (2020a). We also need to choose a kernel that correlates the deviations $\Delta\phi_i$, $\Delta\phi_j$ from ϕ_0 at two values x_i, x_j of the independent variable $x = \ln P$. That is, if the deviation from the fiducial value is drawn from

$$\mathcal{N}(0, \Sigma(\Delta\phi_i, \Delta\phi_j)), \tag{4}$$

where $\mathcal{N}$ represents a normal distribution, then we need to choose the covariance matrix Σ. Following M. C. Miller et al. (2021), we assume that we can write $\Sigma(\Delta\phi_i, \Delta\phi_j) = K(x_i, x_j)$, and we choose a "squared exponential" kernel,

$$K(x, x') = \sigma^2 \exp\left(-\frac{(x - x')^2}{2\ell^2}\right), \tag{5}$$

with $\sigma = 1$ and $\ell = 1$. We note finally that the flexibility of the Gaussian process framework allows the selection of EOSs with $dP/d\epsilon \ll c^2$ in finite density intervals (i.e., effectively, phase transitions; this is the approach taken by R. Essick et al. 2023) or the addition of sharp features in the dependence of sound speed on density that simulate phase transitions and other features (D. Mroczek et al. 2023, 2024).

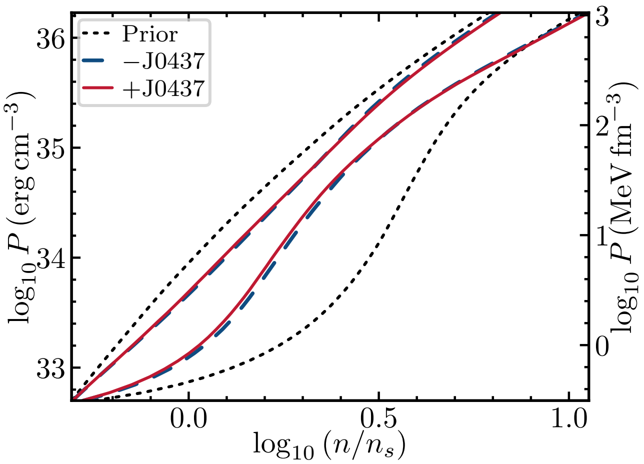

Figure 6. Effect of our mass and radius measurement of PSR J0437−4715 on the EOS of high-density matter, using the Gaussian process framework described in Section 6.2. The black dotted lines show the prior on the EOS, the blue dashed lines show the EOS constraints without our analysis of the NICER data on PSR J0437−4715 (using the measurements discussed in Section 6.1), and the red solid lines show the EOS constraints including our analysis of the PSR J0437−4715 data. Here, we plot the log of the pressure (using different unit systems on the left-hand and on the right-hand axes) versus the log of the baryonic number density in units of the saturation density $n_s = 0.16\ \mathrm{fm}^{-3}$. For each line type, at a given density, the lower curve gives the 5th percentile of the pressure and the upper curve gives the 95th percentile of the pressure. We see that the inclusion of the PSR J0437−4715 measurement tightens the EOS slightly in the $n \approx (1 - 3)n_s$ density range. These results do not include the updated mass of PSR J0952−0607 from R. W. Romani et al. (2026). See the main text for additional details.

6.3. EOS Results

Figure 6 summarizes the effect that our measurement of the mass and radius of PSR J0437−4715 has on our knowledge of the EOS of high-density matter, using the Gaussian process framework described in Section 6.2. In this figure, we compare the 5th and 95th percentiles of the pressure as a function of baryon number density for the prior (black dotted lines), the prior updated with pre-J0437 measurements (blue dashed line; see Section 6.1 for a discussion of the data that we include), and finally the prior additionally updated with the PSR J0437 −4715 measurements of the current Letter (red solid lines). We see that the addition of our PSR J0437−4715 measurement tightens the EOS posterior from $n \approx (1 - 3)n_s$, but only slightly.

Table 4 shows the effect of our updated measurement, as well as the update on the estimated mass of PSR J0348+0432 (from $M = 2.01 \pm 0.04\, M_\odot$ in J. Antoniadis et al. 2013 to $M = 1.806 \pm 0.037\, M_\odot$ in A. Saffer et al. 2025), on the estimate of the maximum mass of a nonrotating neutron star ($M_{\rm TOV}$) and on the radius of a fiducial $M = 1.4\, M_\odot$ neutron star ($R_e(1.4\, M_\odot)$). We see that the effect is small but that both $M_{\rm TOV}$ and $R_e(1.4\, M_\odot)$ are shifted to slightly smaller values. Our results continue to be consistent with $R_e(1.4\, M_\odot) \approx$ 12–13 km, which implies a sharp rise in pressure that is difficult to produce in microscopic models built on standard two- and three-body interactions (S. Reddy 2026, personal communication).

For comparison, in Table 4, we also show the effect on $M_{\rm TOV}$ and $R_e(1.4\, M_\odot)$ if we include the updated mass

Table 4
Maximum Mass and Fiducial Radius Including PSR J0437−4715 Measurement

Quantity	Dataset	-1σ	Median	$+1\sigma$
$M_{\rm TOV}$ ($M_\odot$)	A. J. Dittmann et al. (2024)	2.08	2.22	2.44
$M_{\rm TOV}$ ($M_\odot$)	This work	2.07	2.20	2.44
$M_{\rm TOV}$ ($M_\odot$)	PSR J0952−0607	2.28	2.42	2.61
$R_e(1.4\,M_\odot)$ (km)	A. J. Dittmann et al. (2024)	12.09	12.57	13.06
$R_e(1.4\,M_\odot)$ (km)	This work	12.09	12.56	13.03
$R_e(1.4\,M_\odot)$(km)	PSR J0952−0607	12.06	12.52	12.98

Note. Updated -1σ, median, and $+1\sigma$ points in the posterior distributions of the maximum gravitational mass of a nonrotating neutron star ($M_{\rm TOV}$) and the equatorial circumferential radius of a fiducial nonrotating 1.4 $M_\odot$ neutron star (R_e), compared with those inferred by A. J. Dittmann et al. (2024). In addition to our measurement of PSR J0437−4715 from NICER data, the other update in the "This work" rows involves using the mass estimate for PSR J0348+0432 of $M = 1.806 \pm 0.037\,M_\odot$ from A. Saffer et al. (2025), rather than the older estimate of $M = 2.01 \pm 0.04\,M_\odot$ from J. Antoniadis et al. (2013) that was used in earlier papers. We see that the combination of these two updates shifts both $M_{\rm TOV}$ and $R_e(1.4\,M_\odot)$ to slightly smaller values. For comparison, we also indicate the effect on the $M_{\rm TOV}$ and $R_e(1.4\,M_\odot)$ posteriors if we include the updated mass constraint $M = 2.35 \pm 0.11\,M_\odot$ on the black widow pulsar PSR J0952−0607 from R. W. Romani et al. (2026; see the text for details). We see that the inclusion of this mass would increase the estimated $M_{\rm TOV}$ substantially but would have very little effect on $R_e(1.4\,M_\odot)$.

$M = 2.35 \pm 0.11\,M_\odot$ of the black widow pulsar PSR J0952−0607 from R. W. Romani et al. (2026). A nuance is that because PSR J0952−0607 rotates very rapidly (707.31 Hz), there is extra rotational support. We follow R. W. Romani et al. (2026) in applying a correction of $0.03\,M_\odot$ downward when constraining $M_{\rm TOV}$ (based on the relations presented by A. Konstantinou & S. M. Morsink 2022), so we effectively use a nonrotating mass of $2.32 \pm 0.11\,M_\odot$. We see that the inclusion of this mass would increase the estimated $M_{\rm TOV}$ substantially but would have little effect on $R_e(1.4\,M_\odot)$. We also note that R. W. Romani et al. (2026) urge caution in using the masses of black widows, given possible systematics, but it is clear that if the masses can be considered reliable, then they will provide important input.

7. Conclusion

Models of the NICER data on PSR J0437−4715 require a modulated nonthermal component as well as somewhat complicated spot configurations for the thermal X-ray emission. This complexity offsets the superb data from NICER and the excellent radio measurements from this binary (which provide tight constraints on the mass, the distance, and the observer inclination) to yield a relatively unconstraining radius posterior. This posterior is, however, fully consistent with NICER measurements of the pulsar PSR J0030+0451, which has a similar mass of $M \sim 1.4\,M_\odot$. These measurements, combined with previous analyses of NICER data, the constraints on the tidal deformability of the neutron stars in the gravitational-wave event GW170817, and the high masses inferred for several pulsars, have dramatically improved our understanding of cold, catalyzed matter beyond nuclear saturation density. We anticipate that work in progress on data-efficient filtering and background modeling, together with additional data collection, are likely to result in a modest but significant future improvement in NICER's radius constraint for PSR J0437−4715.

Acknowledgments

We thank the anonymous referee for an unusually constructive and helpful report. We also thank David Blaschke and Sanjay Reddy for helpful suggestions after the initial posting of this Letter. A.J.D. and M.C.M. were supported in part by NASA ADAP grant No. 80NSSC21K0649. Part of this work was performed at the Aspen Center for Physics, which is supported by National Science Foundation grant PHY-2210452. We are grateful to the Institute for Nuclear Theory at the University of Washington for its kind hospitality and stimulating research environment. This research was supported in part by the INT's U.S. Department of Energy grant No. DE-FG02-00ER41132. A.J.D. was supported by NASA through the Hubble Fellowship Program grant No. HST-HF2-51553.001, awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., for NASA, under contract NAS5-26555. J.B. was supported by NASA under award number 80GSFC21M0002. S.B. was supported in part by NASA grants 80NSSC20K0275, 80NSSC22K0728, and 80NSSC25K7544. W.C.G.H. acknowledges support through grant 80NSSC23K0078 from NASA. The research of S.M.M. is supported by NSERC Discovery grant RGPIN-2019-06077. C.C. and Z.W. were supported by NASA under award number 80GSFC24M0006. Some of the resources used in this work were provided by the NASA High-End Computing (HEC) Program through the NASA Center for Climate Simulation (NCCS) at Goddard Space Flight Center. The authors acknowledge the University of Maryland supercomputing resources (http://hpcc.umd.edu) that were made available for conducting the research reported in this Letter. Portions of this work performed at NRL were supported by NASA. This research has made use of data products and software provided by the High Energy Astrophysics Science Archive Research Center (HEASARC), which is a service of the Astrophysics Science Division at NASA/GSFC and the High Energy Astrophysics Division of the Smithsonian Astrophysical Observatory. We acknowledge extensive use of NASA's Astrophysics Data System (ADS) bibliographic services and arXiv.

Facility: NICER (K. C. Gendreau et al. 2016).

Software: emcee (D. Foreman-Mackey et al. 2013), Python (T. E. Oliphant 2007), NumPy (C. R. Harris et al. 2020), Matplotlib (J. D. Hunter 2007), PGF/Ti*k*Z (T. Tantau 2013), Cython (S. Behnel et al. 2011), schwimmbad (A. M. Price-Whelan & D. Foreman-Mackey 2017), HEASoft (Nasa High Energy Astrophysics Science Archive Research Center (Heasarc) 2014), PINT (J. Luo et al. 2021), and pocoMC (M. Karamanis et al. 2022).

Appendix A
Breakdowns of Fits by Components

Figure 7 compares the log evidence and maximum log likelihood for our models. This comparison, like Table 3, demonstrates that, at least for our set of models, a modulated power-law component dramatically improves the fit.

Figure 8 shows the contribution of each of the fit components to the overall best fit, as a function of NICER

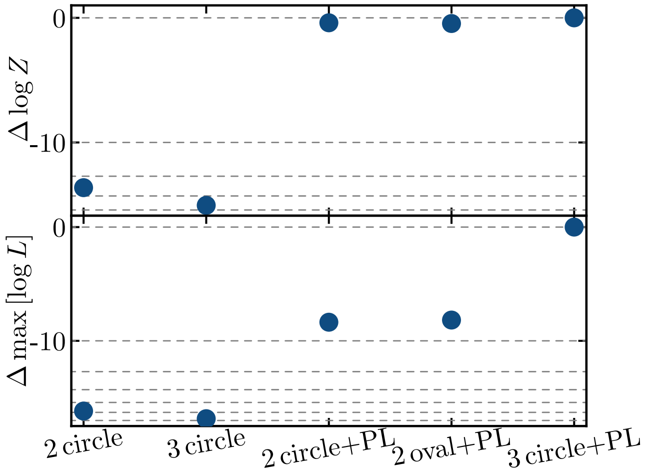

Figure 7. Natural log of evidence (top panel) and natural log of maximum likelihood (bottom panel) for different models, relative to our featured model (three uniform-temperature circular spots plus a modulated Gaussian-profile power law). From the left, these models are numbers 1, 2, 3, 4, and 5 in our list (see Section 4 and Table 2). Models 1 and 2, which do not include a modulated power-law component, are substantially disfavored compared with the models that do include this component. Therefore, to allow those models to fit on the plots, the −10 to 0 portions of both panels are linear, whereas the portions below −10 are logarithmic (the first horizontal dashed line is 10 below the reference model; the second is 20 below; and so on). We conclude that the modulated power-law component must be included in a viable model of this source. We also see that, in terms of the evidence, the two-circle, two-oval, and three-circle models (all with a modulated power law) are all comparable with each other. The three-circle model, however, has a substantially larger maximum log likelihood than the other models.

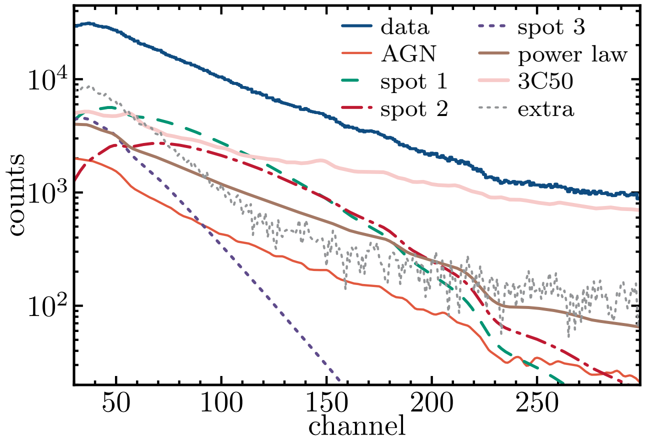

Figure 8. Contribution of each fit component to the phase-summed count spectrum, as a function of NICER PI channel, using our best fit to the data of our model with three uniform-temperature circular spots plus a modulated power law, i.e., our Model 5. The solid blue line shows the data, the fainter solid red line shows the contribution from the empirical NICER background of R. A. Remillard et al. (2022), the brighter solid red line shows the contribution from the angularly nearby AGN, the dashed green line shows the contribution from the first spot, the dotted–dashed red line shows the contribution from the second spot, the bold dotted black line shows the contribution from the third spot, the solid brown line shows the contribution from the modulated power law, and finally the light dotted line shows the remaining contribution, from the extra unmodulated background. Note that the counts axis is logarithmic. This plot demonstrates that multiple components have importance in different channel ranges: the extra (unidentified) unmodulated background is most prominent below channel ∼60, the spots have their greatest relative importance at channels ∼60–120, and although the 3C50 background is the most important single component above channel ∼120, the remaining components all contribute (except for the low-temperature Spot 3). In particular, we see that the power law is the most important of the modulated components above channel ∼200. The apparent erratic nature of the extra component is due to Poisson fluctuations in the data, which are enhanced when other components are subtracted.

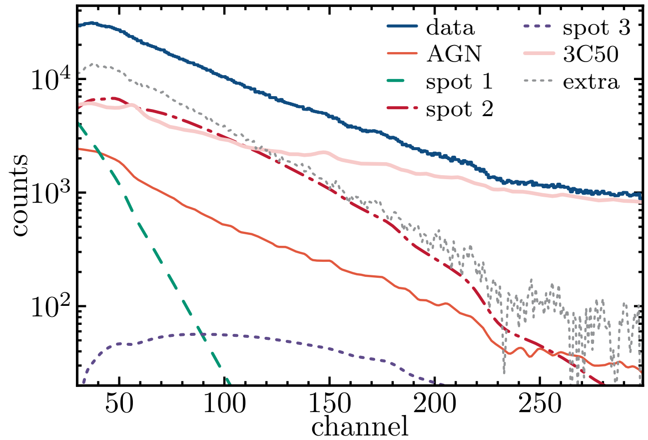

Figure 9. The same as Figure 8 but for the best three-circle fit with no modulated power law. Note that the spots have exchanged roles, e.g., Spot 1 is now the lowest-temperature spot. The overall fit to the spectrum is as good as the fit with a model including a modulated power law, because of the flexibility that we assume in the extra unmodulated background. However, because of the lack of the modulated power-law component, the extra background has more counts.

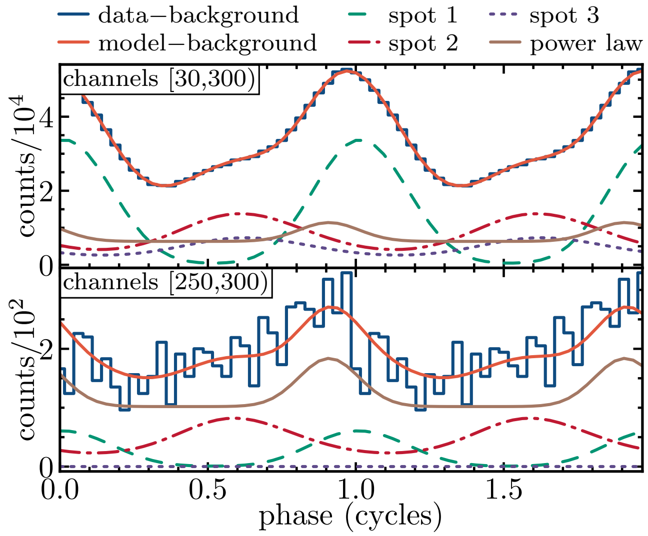

Figure 10. Waveform after subtraction of the unmodulated background components (from the 3C50 background, the angularly nearby AGN, and the "extra" component in Figure 8), using our best fit to the data of our model with three uniform-temperature circular spots plus a modulated power law. The upper panel shows the full bolometric waveform, whereas the lower panel shows the waveform summing NICER PI channels 250 through 299 inclusive. The modulated power law is most important in the higher-energy channels, but it also contributes nonnegligibly to the full bolometric waveform.

PI channel. The upper blue solid line shows the data; the best-fit model, which is the sum of the components, is very close to the data. The three spots are represented by the green dashed line, the red dotted–dashed line, and the blue dotted line. The known background components are shown with the red solid line (from the angularly nearby AGN) and the pink solid line (the empirical background of R. A. Remillard et al. 2022). The brown solid line shows the contribution of the modulated power law, and the gray dotted line shows the additional unmodulated background that we add as part of the fit (note that the vertical axis is logarithmic and that the extra component has picked up the Poisson fluctuations in the data). All of the components are important in at least part of the

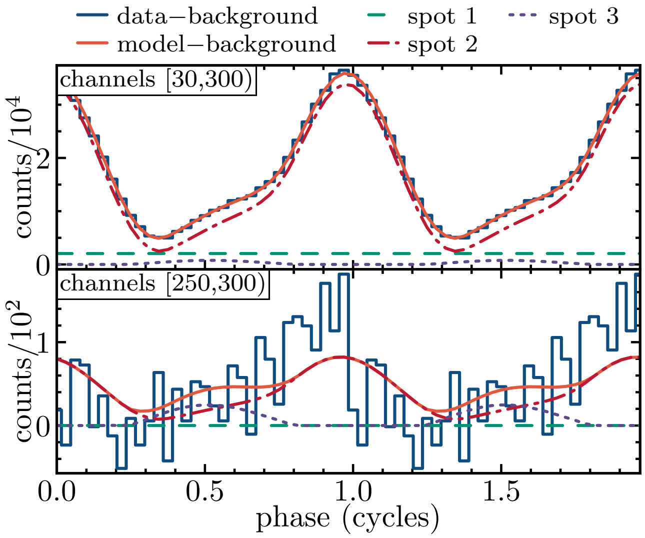

Figure 11. The same as Figure 10 but for the best three-circle fit that does not include a modulated power-law component. Because the spectral contribution from the modulated power law does not exist in this fit, the total background must account for these counts, and thus the data minus background and model minus background count totals are lower than they were in Figure 10. We also see that, without the modulated power law, the fit in the higher-energy channels is considerably worse than it was in Figure 10.

channel range. Figure 9 shows the same plot for the best three-circle model with no modulated power law, where we see that the required extra background is larger than it was when the modulated power law was included.

Figure 10 shows the bolometric (top panel) and high-energy (lower panel) waveforms for the data, the full model, and the model components, after removal of the known and added unmodulated background contributions. This leaves the three spots and the modulated power law. We see that each component contributes; the modulated power law is most important at higher energies but also cannot be neglected in the bolometric waveform. Figure 11 shows the same comparison but for the best three-circle model without a modulated power law.

Appendix B
Comparison of pocoMC and Final Emcee Radius Posterior

Figure 12 shows the progression in the probability distribution for the radius, from the prior (black dotted line) to the final result of the `pocomc` sampling (green dashed line) to the first 10% of the follow-up `emcee` run (red dotted–dashed line; note that the `emcee` run starts from the `pocomc` posterior) to the final 10% of the follow-up `emcee` run (black solid line). The tick marks show the medians of the posterior distributions: 13.89 km, 14.11 km, and 14.15 km, respectively. The `pocomc` posterior is close to, but slightly narrower than, the `emcee` posterior. For comparison, the left-hand panel of Figure 3 in M. C. Miller et al. (2021) shows that, in an analysis of NICER data on PSR J0740+6620 starting from a `MultiNest` run with $N_{\rm live} = 1000$ and a sampling efficiency of SE = 0.01, the median radius increases from ∼12.2 km to ∼13.7 km. Figure 12 shows, in some sense, the worst case for `pocomc`: for the simpler models (Models 1 through 4 in Table 2), the `pocomc` radius posterior is substantially closer to the final `emcee` radius posterior.

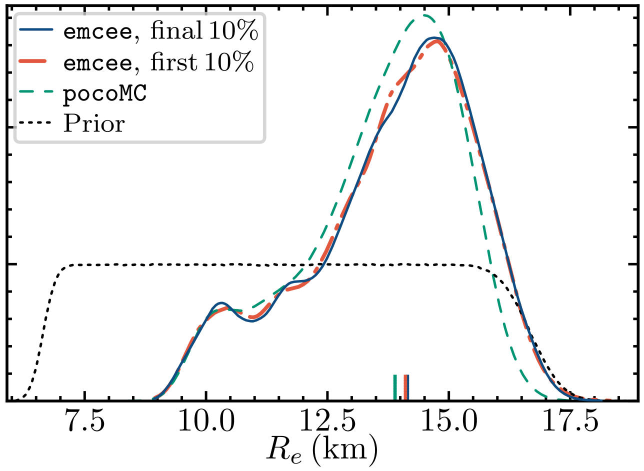

Figure 12. Comparison of several probability distributions for the radius using our model with three uniform-temperature circular spots plus a modulated power law: the prior (dotted black line), the posterior after the completed `pocomc` run (green dashed line), the first 10% of the follow-up `emcee` run (red dotted–dashed line), and the final 10% of the follow-up `emcee` run (solid black line). The vertical axis is linear in the probability density. The colored tick marks rising from the radius axis near $R_e = 14$ km show the medians of each of the probability distributions: 13.89 km for the `pocomc` run, 14.11 km for the first 10% of the `emcee` run, and 14.15 km for the last 10% of the `emcee` run. This figure reinforces the points made in Section 3.8.2: the `pocomc` sampling is broadly consistent with the later `emcee` sampling, although slightly narrower (see also J. U. Lange 2023), and the `emcee` run is well converged.

We note that if the `MultiNest` precision is improved, then the NICER posteriors move closer to the posteriors from samplers such as `emcee` and `pocomc`. For example, T. Salmi et al. (2024a) find in their `MultiNest` analysis of NICER and XMM-Newton data on PSR J0740+6620 a ±1σ equatorial circumferential radius posterior range of $R_{\rm eq} = 12.36^{+1.06}_{-0.80}$ km with $N_{\rm live} = 4096$ and SE = 0.01, but $R_{\rm eq} = 12.55^{+1.37}_{-0.92}$ km with $N_{\rm live} = 4096$ and SE = 0.0001. The latter is substantially closer than the former to the $R_{\rm eq} = 12.76^{+1.49}_{-1.02}$ found by A. Dittmann (2024) using `emcee`. However, for a given precision and accuracy, we have found that `pocomc` has a far lower computational cost than `MultiNest`.

Appendix C
Posterior Distributions

Table 5 lists the median, ±1σ, and ±2σ points in the posterior distributions obtained by fitting our models to only the NICER data, assuming a fully ionized hydrogen atmosphere, and the resulting maximum-likelihood values for each of the parameters in these models. We display the complete corner plot of the posteriors from these same analyses in Figure 13.

Table 5
Fits to NICER Data on PSR J0437−4715

Parameter	Median	-1σ	$+1\sigma$	-2σ	$+2\sigma$	Maximum Likelihood
R_e (km)	14.117	11.880	15.485	9.902	16.541	15.156
GM/c^2R_e	0.149	0.136	0.179	0.128	0.213	0.143
M ($M_\odot$)	1.429	1.385	1.473	1.341	1.517	1.470
$\cos\theta_{c1}$	−0.9916	−0.9931	−0.9883	−0.9941	−0.9799	−0.9928
$\Delta\theta_1$ (rad)	0.580	0.475	0.626	0.316	0.653	0.610
$kT_{\rm eff,1}$ (keV)	0.104	0.100	0.112	0.098	0.123	0.106
$\cos\theta_{c2}$	−0.260	−0.484	0.528	−0.802	0.902	0.847
$\Delta\theta_2$ (rad)	0.011	0.008	0.120	0.006	0.555	0.0097
$kT_{\rm eff,2}$ (keV)	0.166	0.040	0.190	0.021	0.223	0.172
$\Delta\phi_2$ (cycles)	0.618	0.603	0.651	0.592	0.695	0.638
$\cos\theta_{c3}$	0.302	−0.503	0.785	−0.960	0.959	0.925
$\Delta\theta_3$ (rad)	0.218	0.012	1.327	0.007	2.246	0.139
$kT_{\rm eff,3}$ (keV)	0.030	0.018	0.168	0.016	0.207	0.043
$\Delta\phi_3$ (cycles)	0.647	0.613	0.684	0.571	0.772	0.656
$f_{\rm back,3C50}$	0.902	0.874	0.925	0.843	0.946	0.828
$f_{\rm back,AGN}$	0.681	0.545	0.971	0.506	1.435	0.566
$N_{\rm PL}$ (10^{-5} keV^{-1} cm^{-2} s^{-1})	3.080	2.132	4.398	1.465	6.144	6.373
$\alpha_{\rm PL}$	1.670	1.405	1.915	1.130	2.111	2.156
$A_{\rm PL}$	0.861	0.694	0.962	0.517	0.995	0.803
$\psi_{\rm PL}$	0.116	0.075	0.163	0.046	0.217	0.100
$\Delta\psi_{\rm PL}$	0.139	0.094	0.180	0.062	0.247	0.102
$\theta_{\rm obs}$ (rad)	0.7417	0.7414	0.7419	0.7411	0.7422	0.7416
N_H (10^{20} cm^{-2})	0.212	0.070	0.415	0.011	0.700	0.563
d (kpc)	0.1570	0.1568	0.1571	0.1567	0.1573	0.1571
$f_{\rm eff}$	0.999	0.899	1.101	0.798	1.200	0.825

Note. A comparison of the -2σ, -1σ, median, $+1\sigma$, and $+2\sigma$, and maximum-likelihood values inferred from our analysis of NICER data for PSR J0437−4715. Here, we use our featured model, which has three uniform-temperature circular hot spots plus a modulated power-law spectrum with a Gaussian profile in rotational phase.

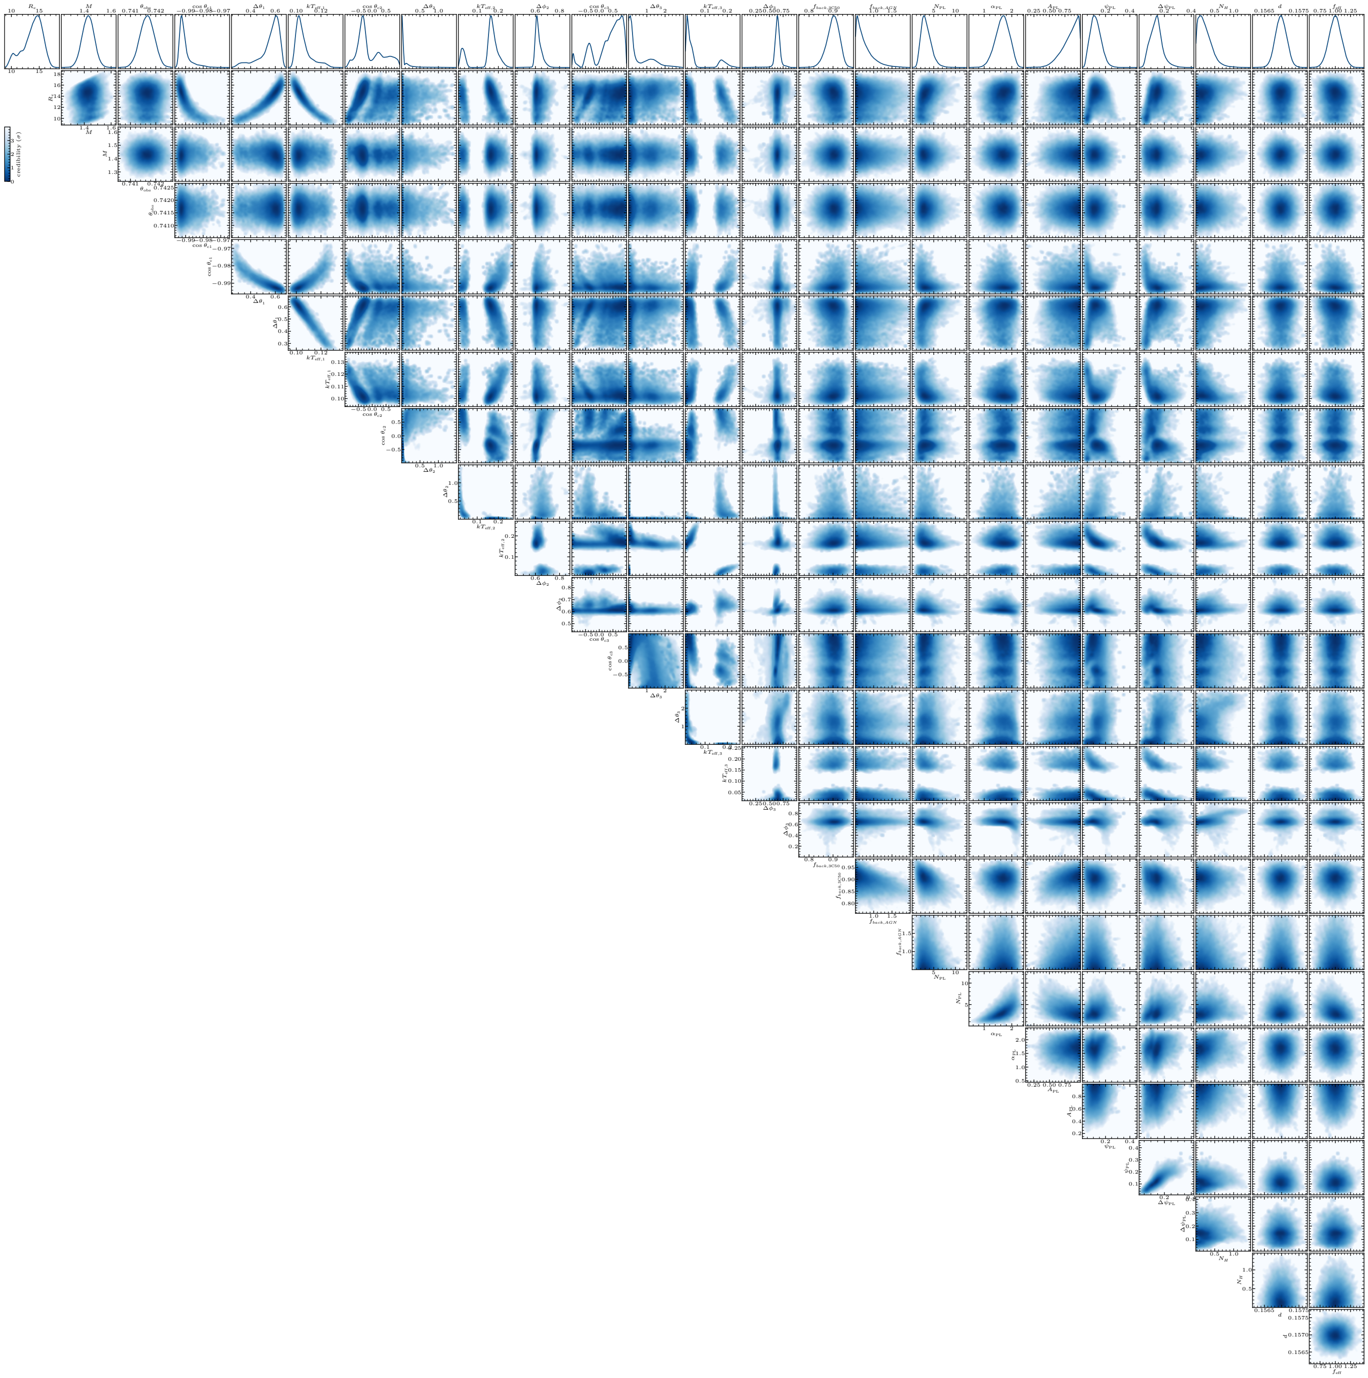

Figure 13. Posterior probability density distributions from our analysis of the NICER data on PSR J0437−4715, where the units are the same as in Table 5. Here, we use our featured model, which has three uniform-temperature circular hot spots plus a modulated power-law spectrum with a Gaussian profile in rotational phase. The power-law normalization $N_{\rm PL}$ is in units of 10^{-5} keV^{-1} cm^{-2} s^{-1}.

ORCID iDs

M. C. Miller https://orcid.org/0000-0002-2666-728X
A. J. Dittmann https://orcid.org/0000-0001-6157-6722
I. M. Holt https://orcid.org/0000-0002-3097-942X
F. K. Lamb https://orcid.org/0000-0002-3862-7402
C. Chirenti https://orcid.org/0000-0003-2759-1368
Z. Arzoumanian https://orcid.org/0009-0008-6187-8753
J. Berteaud https://orcid.org/0000-0003-4962-145X
S. Bogdanov https://orcid.org/0000-0002-9870-2742
K. C. Gendreau https://orcid.org/0000-0001-7115-2819
W. C. G. Ho https://orcid.org/0000-0002-6089-6836
S. M. Morsink https://orcid.org/0000-0003-4357-0575
P. S. Ray https://orcid.org/0000-0002-5297-5278
R. A. Remillard https://orcid.org/0000-0003-4815-0481
Z. Wadiasingh https://orcid.org/0000-0002-9249-0515
M. T. Wolff https://orcid.org/0000-0002-4013-5650

References

Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017, PhRvL, 119, 161101
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2020, ApJL, 892, L3
Agazie, G., Alam, M. F., Anumarlapudi, A., et al. 2023, ApJL, 951, L9
Alcock, C., & Illarionov, A. 1980, ApJ, 235, 534
AlGendy, M., & Morsink, S. M. 2014, ApJ, 791, 78
Antoniadis, J., Freire, P. C. C., Wex, N., et al. 2013, Sci, 340, 448

Bauböck, M., Psaltis, D., & Özel, F. 2019, ApJ, 872, 162
Bauswein, A., Baumgarte, T. W., & Janka, H.-T. 2013, PhRvL, 111, 131101
Baym, G., Furusawa, S., Hatsuda, T., Kojo, T., & Togashi, H. 2019, ApJ, 885, 42
Becker, W., & Trümper, J. 1993, Natur, 365, 528
Becker, W., & Trümper, J. 1999, A&A, 341, 803
Behnel, S., Bradshaw, R., Citro, C., et al. 2011, CSE, 13, 31
Blaschke, D., Alvarez-Castillo, D. E., & Klahn, T. 2016, arXiv:1604.08575
Bogdanov, S. 2013, ApJ, 762, 96
Bogdanov, S., Dittmann, A. J., Ho, W. C. G., et al. 2021, ApJL, 914, L15
Bogdanov, S., Guillot, S., Ray, P. S., et al. 2019a, ApJL, 887, L25
Bogdanov, S., Lamb, F. K., Mahmoodifar, S., et al. 2019b, ApJL, 887, L26
Buchner, J. 2016, Stat. Comput., 26, 383
Buchner, J. 2021, JOSS, 6, 3001
Buchner, J. 2023, Stat. Surv., 17, 169
Chirenti, C., Dichiara, S., Lien, A., Miller, M. C., & Preece, R. 2023, Natur, 613, 253
Chirenti, C., Miller, M. C., Strohmayer, T., & Camp, J. 2019, ApJL, 884, L16
Choudhury, D., Salmi, T., Vinciguerra, S., et al. 2024a, ApJL, 971, L20
Choudhury, D., Watts, A. L., Dittmann, A. J., et al. 2024b, ApJ, 975, 202
Contopoulos, I., & Spitkovsky, A. 2006, ApJ, 643, 1139
Cromartie, H. T., Fonseca, E., Ransom, S. M., et al. 2020, NatAs, 4, 72
De, S., Finstad, D., Lattimer, J. M., et al. 2018, PhRvL, 121, 091102
Demorest, P. B., Pennucci, T., Ransom, S. M., Roberts, M. S. E., & Hessels, J. W. T. 2010, Natur, 467, 1081
Dittmann, A. 2024, OJAp, 7, 79
Dittmann, A. J., Miller, M. C., Lamb, F. K., et al. 2024, ApJ, 974, 295
Drischler, C., Holt, J. W., & Wellenhofer, C. 2021, ARNPS, 71, 403
Essick, R., Landry, P., & Holz, D. E. 2020a, PhRvD, 101, 063007
Essick, R., Legred, I., Chatziioannou, K., Han, S., & Landry, P. 2023, PhRvD, 108, 043013
Essick, R., Tews, I., Landry, P., Reddy, S., & Holz, D. E. 2020b, PhRvC, 102, 055803
Feroz, F., Hobson, M. P., & Bridges, M. 2009, MNRAS, 398, 1601
Finch, E., Legred, I., Chatziioannou, K., et al. 2025, PhRvD, 112, 103023
Fonseca, E., Cromartie, H. T., Pennucci, T. T., et al. 2021, ApJL, 915, L12
Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, PASP, 125, 306
Fryer, C. L., Belczynski, K., Ramirez-Ruiz, E., et al. 2015, ApJ, 812, 24
Gendreau, K. C., Arzoumanian, Z., Adkins, P. W., et al. 2016, SPIE, 9905, 99051H
Goodman, J., & Weare, J. 2010, CAMCS, 5, 65
Gorda, T., Komoltsev, O., & Kurkela, A. 2023, ApJ, 950, 107
Gordon, N., Salmond, D., & Smith, A. 1993, IEE Proceedings F (Radar and Signal Processing), 140, 107
Guedes, V., Radice, D., Chirenti, C., & Yagi, K. 2025, ApJ, 983, 88
Guillot, S., Kaspi, V. M., Archibald, R. F., et al. 2016, MNRAS, 463, 2612
Halpern, J. P., & Marshall, H. L. 1996, ApJ, 464, 760
Handley, W. J., Hobson, M. P., & Lasenby, A. N. 2015, MNRAS, 450, L61
Harding, A. K., & Muslimov, A. G. 2002, ApJ, 568, 862
Harding, A. K., & Muslimov, A. G. 2011, ApJ, 743, 181
Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, Natur, 585, 357
Hebeler, K., Lattimer, J. M., Pethick, C. J., & Schwenk, A. 2013, ApJ, 773, 11
Ho, W. C. G., & Lai, D. 2001, MNRAS, 327, 1081
Holt, I. M., Miller, M. C., Dittmann, A. J., & Lamb, F. K. 2025, arXiv:2511.16759
Hoogkamer, M., Kini, Y., Salmi, T., Watts, A. L., & Buchner, J. 2025, PhRvD, 112, 023008
Hu, H., Kramer, M., Wex, N., Champion, D. J., & Kehl, M. S. 2020, MNRAS, 497, 3118
Hunter, J. D. 2007, CSE, 9, 90
Ih, J., & Kempton, E. M. R. 2021, AJ, 162, 237
Jakab, Z., & Morsink, S. M. 2025, ApJ, 994, 163
Karamanis, M., Beutler, F., Peacock, J. A., Nabergoj, D., & Seljak, U. 2022, MNRAS, 516, 1644
Kitagawa, G. 1996, J. Comput. Graph. Stat., 5, 1
Konstantinou, A., & Morsink, S. M. 2022, ApJ, 934, 139
Kumar, R., Dexheimer, V., Jahan, J., et al. 2024, LRR, 27, 3
Lamb, F. K., Boutloukos, S., Van Wassenhove, S., et al. 2009a, ApJ, 706, 417
Lamb, F. K., Boutloukos, S., Van Wassenhove, S., et al. 2009b, ApJL, 705, L36
Landry, P., & Essick, R. 2019, PhRvD, 99, 084049
Landry, P., Essick, R., & Chatziioannou, K. 2020, PhRvD, 101, 123007
Lange, J. U. 2023, MNRAS, 525, 3181
Lawrence, S., Tervala, J. G., Bedaque, P. F., & Miller, M. C. 2015, ApJ, 808, 186
Legred, I., Chatziioannou, K., Essick, R., Han, S., & Landry, P. 2021, PhRvD, 104, 063003
Legred, I., Sy-Garcia, B. O., Chatziioannou, K., & Essick, R. 2024, PhRvD, 109, 023020
Lemos, P., Weaverdyck, N., Rollins, R. P., et al. 2023, MNRAS, 521, 1184
Li, Z.-X., & Yao, H. 2019, ARCMP, 10, 337
Lindblom, L. 2010, PhRvD, 82, 103011
Lo, K. H., Miller, M. C., Bhattacharyya, S., & Lamb, F. K. 2013, ApJ, 776, 19
Loh, E. Y., J., Gubernatis, J. E., Scalettar, R. T., et al. 1990, PhRvB, 41, 9301
Luo, J., Ransom, S., Demorest, P., et al. 2021, ApJ, 911, 45
Margalit, B., & Metzger, B. D. 2017, ApJL, 850, L19
Mauviard, L., Guillot, S., Salmi, T., et al. 2025, ApJ, 995, 60
Miller, M. C., Chirenti, C., & Lamb, F. K. 2020, ApJ, 888, 12
Miller, M. C., & Lamb, F. K. 2015, ApJ, 808, 31
Miller, M. C., & Lamb, F. K. 2016, EPJA, 52, 63
Miller, M. C., Lamb, F. K., Dittmann, A. J., et al. 2019, ApJL, 887, L24
Miller, M. C., Lamb, F. K., Dittmann, A. J., et al. 2021, ApJL, 918, L28
Morsink, S. M., Leahy, D. A., Cadeau, C., & Braga, J. 2007, ApJ, 663, 1244
Most, E. R., Weih, L. R., Rezzolla, L., & Schaffner-Bielich, J. 2018, PhRvL, 120, 261103
Mroczek, D. 2024, EPJWC, 296, 03002
Mroczek, D., Coleman Miller, M., Noronha-Hostler, J., & Yunes, N. 2023, JPhCS, 2536, 012006
Mroczek, D., Miller, M. C., Noronha-Hostler, J., & Yunes, N. 2024, PhRvD, 110, 123009
Naesseth, C. A., Lindsten, F., & Schön, T. B. 2019, arXiv:1903.04797
Nasa High Energy Astrophysics Science Archive Research Center (Heasarc) 2014, HEAsoft: Unified Release of FTOOLS and XANADU, Astrophysics Source Code Library, ascl: 1408.004
Nättilä, J., Miller, M. C., Steiner, A. W., et al. 2017, A&A, 608, A31
Nelson, B. E., Ford, E. B., Buchner, J., et al. 2020, AJ, 159, 73
Oliphant, T. E. 2007, CSE, 9, 10
Oppenheimer, J. R., & Volkoff, G. M. 1939, PhRv, 55, 374
Özel, F., Psaltis, D., Güver, T., et al. 2016, ApJ, 820, 28
Papamakarios, G., Nalisnick, E., Rezende, D. J., Mohamed, S., & Lakshminarayanan, B. 2021, JMLR, 22, 1, https://jmlr.org/papers/volume22/19-1028/19-1028.pdf
Price-Whelan, A. M., & Foreman-Mackey, D. 2017, JOSS, 2, 357
Rasmussen, C. E., & Williams, C. K. I. 2006, Gaussian Processes for Machine Learning (MIT Press)
Reardon, D. J., Bailes, M., Shannon, R. M., et al. 2024, ApJL, 971, L18
Remillard, R. A., Loewenstein, M., Steiner, J. F., et al. 2022, AJ, 163, 130
Rezzolla, L., Most, E. R., & Weih, L. R. 2018, ApJL, 852, L25
Riley, T. E., Choudhury, D., Salmi, T., et al. 2023, X-PSI: A Python package for neutron star X-ray pulse simulation and inference, v1.2.2, Zenodo, doi:10.5281/zenodo.7632629
Riley, T. E., Watts, A. L., Bogdanov, S., et al. 2019, ApJL, 887, L21
Riley, T. E., Watts, A. L., Ray, P. S., et al. 2021, ApJL, 918, L27
Romani, R. W., Beleznay, M., Filippenko, A. V., Brink, T. G., & Zheng, W. 2026, ApJ, 996, 101
Romani, R. W., Kandel, D., Filippenko, A. V., Brink, T. G., & Zheng, W. 2022, ApJL, 934, L17
Ruiz, M., Shapiro, S. L., & Tsokaros, A. 2018, PhRvD, 97, 021501
Saffer, A., Fonseca, E., Ransom, S., et al. 2025, ApJL, 983, L20
Salmi, T., Choudhury, D., Kini, Y., et al. 2024a, ApJ, 974, 294
Salmi, T., Deneva, J. S., Ray, P. S., et al. 2024b, ApJ, 976, 58
Salmi, T., Suleimanov, V. F., Nättilä, J., & Poutanen, J. 2020, A&A, 641, A15
Salmi, T., Vinciguerra, S., Choudhury, D., et al. 2022, ApJ, 941, 150
Shklovskii, I. S. 1970, SvA, 13, 562
Silverman, B. W. 1986, Density Estimation for Statistics and Data Analysis (Chapman & Hall/CRC)
Skilling, J. 2004, AIPC, 735, 395
Speagle, J. S. 2020, MNRAS, 493, 3132
Stewart, L., & McCarty, P., Jr 1992, SPIE, 1699, 177
Tantau, T. 2013, The TikZ and PGF Packages, SourceForge, http://sourceforge.net/projects/pgf/
Tolman, R. C. 1939, PhRv, 55, 364
Tsang, M. B., Stone, J. R., Camera, F., et al. 2012, PhRvC, 86, 015803
Verbiest, J. P. W., Bailes, M., van Straten, W., et al. 2008, ApJ, 679, 675
Vinciguerra, S., Salmi, T., Watts, A. L., et al. 2024, ApJ, 961, 62
Watts, A. L., Andersson, N., Chakrabarty, D., et al. 2016, RvMP, 88, 021001
Weinberg, S. 1990, PhLB, 251, 288
Weinberg, S. 1991, NuPhB, 363, 3
Weinberg, S. 1992, PhLB, 295, 114
Williams, M. J., Karamanis, M., Luo, Y., & Seljak, U. 2025, MNRAS, 543, 1479
Zhao, T., Psaltis, D., & Özel, F. 2025, ApJ, 982, 112