



Constraining the Neutron Star Mass–Radius Relation and Dense Matter Equation of State with NICER. III. Model Description and Verification of Parameter Estimation Codes

Slavko Bogdanov¹, Alexander J. Dittmann^{2,3}, Wynn C. G. Ho⁴, Frederick K. Lamb^{5,6}, Simin Mahmoodifar⁷, M. Coleman Miller², Sharon M. Morsink⁸, Thomas E. Riley⁹, Tod E. Strohmayer⁷, Anna L. Watts⁹, Devarshi Choudhury⁹, Sebastien Guillot^{10,11}, Alice K. Harding¹², Paul S. Ray¹³, Zorawar Wadiasingh^{14,15}, Michael T. Wolff¹³, Craig B. Markwardt¹⁶, Zaven Arzoumanian¹⁶, and Keith C. Gendreau¹⁶

¹ Columbia Astrophysics Laboratory, Columbia University, 550 West 120th Street, New York, NY 10027, USA; slavko@astro.columbia.edu

² Department of Astronomy and Joint Space-Science Institute, University of Maryland, College Park, MD 20742-2421, USA

³ Center for Computational Astrophysics, Flatiron Institute, 162 5th Avenue, New York, NY 10010, USA

⁴ Department of Physics and Astronomy, Haverford College, 370 Lancaster Avenue, Haverford, PA 19041, USA

⁵ 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

⁶ Department of Astronomy, University of Illinois at Urbana-Champaign, 1002 West Green Street, Urbana, IL 61801-3074, USA

⁷ Astrophysics Science Division and Joint Space-Science Institute, NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA

⁸ Department of Physics, University of Alberta, Edmonton, AB T6G 2G7, Canada

⁹ Anton Pannekoek Institute for Astronomy, University of Amsterdam, Science Park 904, 1090GE Amsterdam, The Netherlands

¹⁰ IRAP, CNRS, 9 avenue du Colonel Roche, BP 44346, F-31028 Toulouse Cedex 4, France

¹¹ Université de Toulouse, CNES, UPS-OMP, F-31028 Toulouse, France

¹² Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM 58545, USA

¹³ Space Science Division, U.S. Naval Research Laboratory, Washington, DC 20375, USA

¹⁴ Astrophysics Science Division, NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA

¹⁵ Universities Space Research Association (USRA), Columbia, MD 21046, USA

¹⁶ X-Ray Astrophysics Laboratory, NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA

Received 2020 November 23; revised 2021 March 26; accepted 2021 April 15; published 2021 June 9

Abstract

We describe the X-ray pulse profile models we use and how we use them to analyze Neutron Star Interior Composition Explorer (NICER) observations of rotation-powered millisecond pulsars to obtain information about the mass–radius relation of neutron stars and the equation of state of the dense matter in their cores. Here we detail our modeling of the observed profile of PSR J0030+0451 that we analyzed in Miller et al. and Riley et al. and describe a cross-verification of computations of the pulse profiles of a star with $R/M \approx 3$, in case stars this compact need to be considered in future analyses. We also present our early cross-verification efforts of the parameter estimation procedures used by Miller et al. and Riley et al. by analyzing two distinct synthetic data sets. Both codes yielded credible regions in the mass–radius plane that are statistically consistent with one another, and both gave posterior distributions for model parameter values consistent with the values that were used to generate the data. We also summarize the additional tests of the parameter estimation procedure of Miller et al. that used synthetic pulse profiles and the NICER pulse profile of PSR J0030+0451. We then illustrate how the precision of mass and radius estimates depends on the pulsar’s spin rate and the size of its hot spot by analyzing four different synthetic pulse profiles. Finally, we assess possible sources of systematic error in the estimates made using this technique, some of which may warrant further investigation.

Unified Astronomy Thesaurus concepts: Neutron stars (1108); Compact objects (288); Nuclear astrophysics (1129); Pulsars (1306); Neutron star cores (1107); X-ray astronomy (1810); X-ray sources (1822); Diffuse x-ray background (384); Stellar atmospheres (1584)

1. Introduction

Strongly degenerate matter at densities above the density of nuclear saturation ($\rho_s \equiv 2.4 \times 10^{14} \text{ g cm}^{-3}$) can exist stably in the cores of neutron stars (NSs). This makes NSs important natural laboratories for studying the physics of the strong interaction and the state of matter at supranuclear densities, especially because at present, experimental and theoretical approaches to understanding dense matter suffer from multiple challenges (see, e.g., Watts et al. 2016 for a review). A number of theoretically well-motivated proposals for the composition and properties of matter in the low-temperature, high-density regime are consistent with current data; they range from mainly nucleonic matter to matter with hyperons, deconfined quarks, color superconducting phases, or Bose–Einstein condensates

(see, e.g., Lattimer & Prakash 2016; Oertel et al. 2017; Baym et al. 2018). Because the mass–radius relation of NSs depends sensitively on the pressure–density relation of the matter in their interiors (see, e.g., Lattimer & Prakash 2001, 2005; Özel & Psaltis 2009; Read et al. 2009; Hebeler et al. 2013; Steiner et al. 2013), measurements of the mass and radius of several NSs to a precision of a few percent can, in principle, provide insights into the composition and properties of the matter in their cores, especially if the stars have different masses. This has prompted the development of various astrophysical techniques for inferring the masses and radii of NSs. Several such methods use observations of thermal electromagnetic radiation from the surface of the NS, which falls in the X-ray band when the temperature of the surface is $\gtrsim 0.1 \text{ keV}$ (see, e.g., Heinke 2013; Miller 2013; Özel 2013; Potekhin 2014;

Miller & Lamb 2016; and references therein). Detailed analyses of the pulse profiles of the pulsed X-ray emission produced by “hot spots” on the surfaces of rotating NSs have been identified as a particularly promising approach for determining the mass–radius relation of NSs. This technique has been studied extensively over the past four decades (Pechenick et al. 1983; Strohmayer 1992; Page 1995; Pavlov & Zavlin 1997; Miller & Lamb 1998, 2015; Braje et al. 2000; Weinberg et al. 2001; Beloborodov 2002; Poutanen & Gierliński 2003; Cadeau et al. 2007; Morsink et al. 2007; Lo et al. 2013; AlGendy & Morsink 2014; Psaltis & Özel 2014; Näätä & Pihajoki 2018).

One of the primary scientific goals of the Neutron Star Interior Composition Explorer (NICER) mission (Gendreau et al. 2016) is to make precise and reliable measurements of the radii and masses of several NSs. To accomplish this, NICER is conducting extensive observations of a promising set of nearby rotation-powered millisecond pulsars (MSPs) that exhibit thermal X-ray pulsations. This paper is the third in a series of methodological papers describing how this approach can be used to measure the NS mass–radius relation and the dense matter equation of state (EoS) using NICER. Bogdanov et al. (2019a, hereafter Paper I) described the NICER observations and the resulting MSP data that are being collected, while Bogdanov et al. (2019b, hereafter Paper II) described the approach and codes used to compute the pulse profile models that are being compared with the data being collected by NICER. The purpose of the present work (Paper III) is (i) to provide a compilation of the details of the pulse profile modeling technique that was applied to the NICER data on PSR J0030+0451 by Miller et al. (2019) and Riley et al. (2019) and will be used in future analyses of pulse profile data obtained with NICER; (ii) to describe the tests and cross-comparisons we have performed to check the parameter estimation procedures that were used in Miller et al. (2019) and Riley et al. (2019) and will be employed in future work; (iii) to provide an illustrative comparison between the mass–radius constraints that can be obtained for different assumed spins and hot-spot sizes; and (iv) to provide a brief assessment of the possible systematic errors in this approach. The mass–radius posterior distributions for PSR J0030+0451 that were presented in Miller et al. (2019) and Riley et al. (2019) were computed using these codes and the pulse profile models described here. Miller et al. (2019) and Raaijmakers et al. (2019, 2020) then used these posteriors to obtain constraints on the EoS of cold, dense matter. The paper is organized as follows. In Section 2, we detail our modeling of the X-ray emission from the surface of a rapidly rotating NS and the resulting energy-resolved X-ray pulse profile models we use to constrain the mass–radius relation of NSs and the EoS of the dense matter in their interiors. This section also reports a comparison of two independent computations of the pulse profile produced by a star with $R_{\text{eq}}/M \approx 3$, performed to cross-check the accuracy of these codes in case stars this compact need to be considered in future analyses of NICER pulse profile data. In Section 3, we describe the tests and cross-comparisons we have performed to verify the two different parameter estimation procedures presented in Miller et al. (2019) and Riley et al. (2019). The first two tests we describe in detail here were carried out early in our development of parameter estimation procedures. This section also summarizes the results of two additional tests of the parameter estimation procedure used by Miller et al. (2019), using a synthetic pulse profile that

was designed to mimic the pulse profile of PSR J0030+0451 observed with NICER. In Section 4, we demonstrate how the NS’s spin rate and the size of its hot spot affect the constraints on the star’s mass and radius that can be obtained using the pulse profile fitting method using synthetic data. In Section 5, we briefly discuss and assess several possible sources of systematic error in estimates of stellar masses and radii made using this technique and evaluate whether these warrant further investigation. We summarize the contents and discussion in this paper in Section 6.

2. Model Components

There are numerous physical processes that determine the characteristics of an X-ray pulse profile observed far from an X-ray-emitting, rotating NS. These include the generation of the X-rays that emerge from the star’s surface, the propagation of this radiation in the exterior spacetime of the NS and then through the interstellar medium (ISM) between the star and Earth, and the interaction of the X-rays with the optics and detectors of the observing telescope. In more detail, the effects that must be considered are as follows.

1. The physical properties of the emitting atmosphere of the star (including its density, chemical composition, ionization state, and magnetization) and the spectrum and beaming pattern of the radiation that emerges from each point on the star.
2. The temperature distribution on the stellar surface, which we model as spots with various locations, sizes, shapes, and temperatures, each of which emits radiation that contributes to the total flux potentially observed using NICER as a function of the star’s rotational phase.
3. The special relativistic boost and aberration of the emitted radiation caused by the rotational velocity of the emitting gas.
4. The general relativistic gravitational lensing of the radiation emitted by the NS surface.
5. The partial or complete occultation of each hot spot by the body of the star as it rotates.
6. The general relativistic redshift and time dilation that affects the frequency and travel time of the radiation as it propagates from the emitting surface to the observer.
7. The energy-dependent attenuation of the emitted X-rays by the ISM between the star and the observer.
8. The propagation of the incident X-ray photons through the NICER concentrator optics and the efficiency of the detectors in registering them, both as functions of energy, information that is encapsulated in the redistribution matrix file (RMF) and the ancillary response file (ARF).
9. The combined backgrounds produced by emission from the nonspot portions of the star, other parts of the pulsar system (including any wind nebula and/or binary companion), neighboring astronomical sources, and any diffuse emission in the NICER field of view, as well as the background produced by the radiation environment of the telescope.

We now describe each of these components of the model in more detail, including how they are incorporated into our analyses of NICER data on MSPs. Items 3–6 were discussed in depth in Paper II.




Figure 1. Polar diagrams of the specific intensity I_ν produced by partially ionized hydrogen (solid), fully ionized hydrogen (dotted), partially ionized helium (long-dashed), and fully ionized helium (short-dashed) unmagnetized atmosphere models at the four photon energies indicated (0.5, 0.8, 1.0, and 2.0 keV), computed using the NSX code assuming $T_{\text{eff}} = 10^6$ K and $g_s = 2.4 \times 10^{14}$ cm s $^{-2}$.

2.1. NS Model Atmospheres

The thermal X-rays observed from rotation-powered MSPs most likely originate in heated regions of a geometrically thin but optically thick gaseous atmosphere that covers the stellar surface. The properties of this atmosphere, such as the strength of the magnetic field within it, its chemical composition, and its ionization state, directly determine the spectrum and beaming pattern of the radiation that emerges from the atmosphere (see, e.g., Potekhin 2014, for a review).

Magnetic fields—The characteristic strength at which the magnetic field in an NS atmosphere has a large effect on the radiation that emerges from it is $B_0 \equiv e^3 m_e^2 c / \hbar^3 = 2 \times 10^9$ G (see Lai 2001 for a review). At field strengths at or in excess of B_0 , the electrons in the atoms and molecules in the atmosphere are more tightly bound to their nuclei than in weaker fields, their emission is strongly anisotropic and polarized (see, e.g., Mészáros 1992), and spectral features produced by the electron cyclotron resonance and electron bound states may be visible in the spectrum of the radiation that emerges from the atmosphere (see, e.g., Lai 2001; Potekhin 2010, 2014; Suleimanov et al. 2010, 2012).

The external magnetic fields of rotation-powered pulsars have most commonly been assumed to be dipolar, with the dipole moment at the center of the star. With this assumption, the pulsar’s magnetic dipole moment can be estimated from its

pulse period P and period derivative \dot{P} using the braking formula for a force-free dipolar magnetic field¹⁷ (see, e.g., Contopoulos et al. 1999; Contopoulos & Spitkovsky 2006; Spitkovsky 2006). For some of the MSPs of interest to us, the strength of the surface magnetic field inferred by making these assumptions is $\sim 5 \times 10^8$ G and hence significantly smaller than B_0 (see Miller et al. 2019). Thus, if the magnetic fields of these MSPs were centered dipoles, the appropriate force-free braking formula would indicate surface fields that would have only a small effect on the radiation produced there, and spectra and beaming patterns computed using nonmagnetic model atmospheres (such as those shown in Figure 1) would be adequate.

The preferred configuration of the temperature distribution for PSR J0030+0451 found by the independent analyses of Miller et al. (2019) and Riley et al. (2019) consists of one highly elongated (in the longitudinal direction) region and a second smaller, less elongated region, both on the hemisphere below the stellar equator relative to the observer’s line of sight. Such a temperature distribution is difficult to reconcile with a standard dipole field but could be reproduced by assuming a dipole+quadrupole field¹⁸ (see, e.g., Gralla et al. 2017; Lockhart et al. 2019; Chen et al. 2020; Kalapotharakos et al. 2021). In such a case, the strength of the surface magnetic field inferred using a magnetic dipole braking formula could

¹⁷ Studies of the radio and γ -ray emission of γ -ray pulsars have provided increasing evidence that their magnetic field structure resembles that predicted by force-free models; see Kalapotharakos et al. (2019) and references therein.

¹⁸ The implications of such a field geometry for different aspects of pulsars are discussed in Bilous et al. (2019) and Miller et al. (2019).

substantially underestimate the strength of the magnetic field within the atmospheres of hot spots (see Miller et al. 2019). Even if the star’s external magnetic field is predominantly dipolar, if it is off-center, the dipole component of the magnetic field at the stellar surface could be ~ 5 –10 times larger than would be inferred from the centered-dipole braking formula if the same magnetic field were centered in the star. In that case, the strength of the surface magnetic field near the magnetic poles could be larger than B_0 . The absence of any obvious narrowband features such as absorption lines in the spectrum of PSR J0030+0415 observed using XMM-Newton (see, e.g., Bogdanov & Grindlay 2009) and the ~ 0.3 –3 keV spectrum observed using NICER argues against surface magnetic fields with strengths $B \approx (3\text{--}30) \times 10^{10}$ G (assuming a gravitational redshift at the stellar surface of 20%). While electron cyclotron features produced by much stronger magnetic fields $B \gg 10^{11}$ G would appear at $E > 3$ keV and therefore could be missed by these observations, such strong magnetic fields would create pencil or fan-beam emission, which would probably produce markedly different pulse profiles than the X-ray pulse profile observed using NICER. Comparisons between beaming patterns produced by nonmagnetic and magnetic hydrogen atmospheres with an effective temperature of 10^6 K show differences much greater than the differences between hydrogen and helium (see Figure 1).

As discussed by Miller et al. (2019) and Riley et al. (2019), energy-resolved pulse profile models constructed by assuming negligible magnetic fields in the atmospheres of the hot spots are consistent with the energy-resolved pulse profiles observed using NICER. At present, there is no evidence that the available data require that the strengths of the magnetic fields in the atmospheres of the hot spots approach or exceed B_0 , but the possibility of subtle effects introduced by strong fields may need to be examined (see Section 5).

Chemical composition—During their accretion-powered spin-up phases, the chemical composition of the atmospheres of MSPs is expected to be determined primarily by the composition of the gas transferred from the envelope of the companion star. In many cases, the transferred gas is expected to be primarily hydrogen. Even if helium or heavier elements dominate in the accreted gas, some hydrogen is expected to be created during the accretion phase by spallation (Bildsten et al. 1992; Randhawa et al. 2019).

The rotation-powered MSPs considered here are not expected to be actively accreting. Once accretion has ended, the heavier elements in their outer layers are expected to sink rapidly via diffusive gravitational separation, leaving behind the lightest element (Alcock & Illarionov 1980; Hameury et al. 1983; Brown et al. 2002). The upper atmospheres of rotation-powered MSPs are therefore expected to consist of the lightest element that has been accreted, which in most cases is expected to be hydrogen. Accretion of even a small amount of hydrogen after mass transfer from the companion has ended would also be sufficient to create an upper atmosphere of pure hydrogen (Blaes et al. 1992; Wijngaarden et al. 2019, 2020).

The upper atmosphere of a rotation-powered MSP could be helium if the hydrogen originally present in its atmosphere has all been converted to helium via diffusive nuclear burning (Rosen 1968; Chang & Bildsten 2003, 2004; Wijngaarden et al. 2019) or it accreted pure helium from a helium white dwarf companion during its mass-transfer phase. An MSP could have an atmosphere with even heavier elements, if no hydrogen or

helium was ever accreted, if all of the accreted hydrogen and helium has been burned to heavier elements by nuclear reactions or it has had a wind that has excavated its surface to expose underlying heavy elements (Chang & Bildsten 2004). A comparison between beaming patterns produced by a helium atmosphere and a hydrogen atmosphere with an effective temperature of 10^6 K shows that they differ by $< 5\%$ at 0.5–1 keV and $< 30\%$ at 2 keV (see Figure 1).

Ionization state—The opacity tables that are currently available for partially ionized hydrogen and helium (Iglesias & Rogers 1996; Badnell et al. 2005; Colgan et al. 2016) do not cover the entire range of energies and temperatures needed to model the pulse profiles of the MSPs we consider. However, the neutral fraction is low in atmospheres with the inferred temperatures (≈ 0.5 – 2×10^6 K) of the hot spots in these MSPs, and a comparison between the beaming patterns produced by a partially ionized atmosphere constructed using OP¹⁹ opacities (Badnell et al. 2005) and a fully ionized atmosphere with an effective temperature of 10^6 K shows that they differ by $< 2\%$ at 0.5–1 keV and $< 11\%$ at 2 keV (see Figure 1). We have therefore used fully ionized model atmospheres in our work to date.

To model the pulse profiles and phase-resolved spectra of MSPs, one needs to determine the specific intensity $I_\nu(\mu)$ from a local patch of the NS surface, where ν is the frequency of the radiation and $\mu = \cos \theta$ in terms of the angle θ between the direction of propagation of the radiation and the local normal to the surface. Here $I_\nu(\mu)$ is computed by solving the radiative transfer equation (RTE) using established techniques (Mihalas 1978). Early efforts to compute realistic spectra considered emission from (nonmagnetic) fully ionized, light-element atmospheres and anisotropic beaming patterns that depended on the chemical composition of the atmosphere and the photon energy (Romani 1987; Rajagopal & Romani 1996; Zavlin et al. 1996). Since then, many nonmagnetic NS model atmospheres have been constructed (Gänssicke et al. 2002; McClintock et al. 2004; Heinke et al. 2006; Ho & Heinke 2009; Haakonsen et al. 2012; Suleimanov et al. 2014; González-Caniulef et al. 2019). For the parameter estimation analyses of the NICER data presented in Miller et al. (2019) and Riley et al. (2019), we used the NSX fully ionized H atmosphere models of Ho & Lai (2001) and Ho & Heinke (2009). For comparison, we also computed fully ionized He atmosphere models. We plan to use these models in future analyses to explore the sensitivity of the inferred M and R credible regions to the assumed chemical composition, among other things. In Section 5, we discuss the possible systematic errors that may be introduced by our assumption of a hydrogen atmosphere with a negligible magnetic field. We note that non-local thermodynamic equilibrium (NLTE) effects (see, e.g., Rauch et al. 2008) within the atmosphere are not important at temperatures and energies relevant for rotation-powered MSPs ($\sim 10^6$ K and $\lesssim 1$ keV). The same holds true for Compton scattering, which only becomes important at temperatures $\gtrsim 3 \times 10^6$ K; deviations at lower temperatures only start to become noticeable at photon energies $E > 2$ keV but are still much less than 1% and therefore would have no significant impact on our analysis (Suleimanov & Werner 2007; Salmi et al. 2019), given that nearly all of the thermal emission observed from the MSPs observed by NICER occurs below ~ 2 keV.

¹⁹ The Opacity Project. See <http://cdsweb.u-strasbg.fr/topbase/OpacityTables.html>.

Computational method—Determining the specific intensity $I(E, \theta)$ emerging from an atmosphere as a function of the photon energy E and emission angle θ is computationally expensive. For this reason, in our analyses of the NICER pulse profile data, we use precomputed, high-resolution lookup tables of $I(E, \theta)$ for grids in the logarithm of the effective temperature, $\log T_{\text{eff}}$; the logarithm of the surface gravity, $\log g_s$; and the emission angle with respect to the surface normal θ , expressed as $\mu = \cos \theta$ (see Ho & Lai 2001 for details). For practical reasons, we specified the energy grid in terms of $\log(E/kT_{\text{eff}})$ and the specific intensity grid in terms of $\log(I/T_{\text{eff}}^3)$, because E/kT_{eff} and I/T_{eff}^3 are relativistically invariant quantities. The ranges of $\log T_{\text{eff}}$ covered in our tables were [5.1, 6.8] for the hydrogen atmosphere and [5.1, 6.5] for the helium atmosphere. The range for $\log g_s$ was [13.7, 14.7], the range for μ was [1×10^{-6} , $1 - 1 \times 10^{-6}$] (plus three additional values near $\theta = 0$ and $\pi/2$ to improve the grid resolution near and avoid singularities at these values of μ), and the range for $\log(E/kT_{\text{eff}})$ was [−1.3, 2.0]. Based on tests, we found that spacings of the entries within these ranges of $\Delta(\log T_{\text{eff}}) = 0.05$, $\Delta(\log g_s) = 0.1$, $\Delta\mu = 1/60$, and $\Delta \log(E/kT_{\text{eff}}) = 0.0338$ produced interpolation errors that were negligible for our purposes ($\lesssim 0.1\%$). To obtain the specific intensity for a particular observed photon energy E from a surface element with an effective temperature T_{eff} , we took the corresponding value of $\log(E/kT_{\text{eff}})$ and (1) computed the surface gravity g_s at the colatitude θ_c of the surface element in question, using the expression for the surface gravity given by AlGendy & Morsink (2014) for the assumed combination of NS mass, radius, and rotational frequency; (2) determined the value of μ for which a photon from that point on the surface reaches the observer (for the given combination of surface element colatitude and rotational phase ϕ ; see Paper II for details); (3) identified the entries in the lookup table that bracket the required values of $\log(E/kT_{\text{eff}})$, $\log T_{\text{eff}}$, $\log g_s$, and μ ; and (4) performed quadratic or quartic polynomial (Miller et al. 2019) or cubic polynomial (Riley et al. 2019) interpolation to obtain the corresponding value of $\log(I/T_{\text{eff}}^3)$.

While the physics behind light-element nonmagnetic NS atmospheres is well understood, we caution that there may exist differences in the details of the numerical implementation, such as the discretization scheme, the depth considered for the atmosphere, the criterion for convergence, etc. With this in mind, we compared the output of NSX with the publicly available H atmosphere model McPHAC²⁰ (Haakonsen et al. 2012). In comparing NSX with McPHAC, we identified two major weaknesses of McPHAC that greatly limit its convergence to a consistent solution of the RTE. First, the angular grid and integration make use of an even number of Gauss–Legendre abscissas. This leads to a lowest absolute value of the abscissa that differs significantly from zero, unless a very large number of grid points is used, and as a result of this nonzero value, McPHAC runs slowly. This problem can be rectified by a trivial modification of McPHAC, so that it computes an odd number of abscissas for the angular grid and allows the lowest value to be set arbitrarily close to zero. We find that such a modification yields a greater-than-tenfold improvement in computational speed for the cases we consider. Second, the temperature-correction scheme used by McPHAC to converge to the solution of the RTE has several defects, as described in

Mihalas (1978). We further note that in the publicly available version of McPHAC, the physical constants used are either outdated or defined with low precision. In our comparisons against NSX and the single hot-spot test described in Section 3.1, we modified the values of these constants to their current and full precision values. In the two hot-spot tests described in Section 3.2, we used the NSX nonmagnetic, fully ionized hydrogen atmosphere model.

2.2. Modeling the Hot Spots

In the parameter estimation codes used by Miller et al. (2019) and Riley et al. (2019) to model the NICER data, the X-ray emitting areas on the stellar surface were represented as a combination of single-temperature regions (see both papers for more in-depth descriptions of their modeling procedures, choices of priors, and treatment of degeneracies).

In Riley et al. (2019), each single-temperature region is represented geometrically by a spherical cap with a size specified by an angular radius, with regions allowed to overlap in different ways. If a hot region is constituted by one such single-temperature region, it has the morphology of a hot spot. More complex hot regions can be constituted by two single-temperature regions that partially or wholly overlap, where one takes precedence. An order of precedence means that when evaluating the radiation intensity at a spacetime event (on the surface) that belongs to the intersection of two single-temperature regions, one of the regions supersedes the other, meaning one temperature is ignored at that spacetime event. Moreover, one of the single-temperature regions can have a temperature so low that the X-ray signal it generates is not computed because it would contribute negligibly in the NICER wave band; however, if that (effectively) nonradiating region also takes precedence, the overall hot region is not a simple spherical cap. This implementation permits the representation of a variety of hot region geometries, for example, symmetric or asymmetric ring and crescent shapes²¹ and two-temperature configurations where the hotter component is itself ringlike, a crescent, or simply a spherical cap. Finally, when two disjoint regions on the star are considered, these hot regions are not permitted to overlap (which would require extension of the order of precedence); this condition is used to define the prior support of hot region configurations. For the PSR J0030+0451 analysis in Riley et al. (2019), one of the X-ray emitting areas was found to be best represented by a hot crescent (a cooler spherical-cap component was shown not to be useful).

In their modeling of soft X-ray pulse profiles, Miller et al. (2019) started with the simplest possible description of the heated region(s) on the stellar surface and then, as necessary, considered more complicated descriptions, guided at each step by the differences between the best-fit simpler pulse profile models and the pulsations observed using NICER. Their algorithm allows circular and oval hot spots that can overlap in any way that improves the fit. In the case of PSR J0030+0415, they began by considering models with two and then three and four different uniform-temperature circular hot spots and were led to models with two and three different uniform-temperature oval spots. As the fitting proceeds, their algorithm allows the

²⁰ Available for download at <https://github.com/McPHAC/>.

²¹ A ring can be constructed if the single-temperature region that takes precedence has a low enough temperature so as not to contribute significantly to the observed emission and is itself a proper subset of the other single-temperature region. Similarly, a crescent is constituted by two such regions where the overlap is only partial. See Riley et al. (2019) for further details.

heated regions in the model to evolve toward a variety of different complex shapes, including multiple separate hot spots, multiple-temperature spots in which each spot has hotter and cooler regions, elongated or more complicated configurations produced by two or more adjacent or partially or completely overlapping circular or oval hot spots, as well as crescent-shaped hot regions produced by dark circular or oval spots that partially or completely overlap hot circular or oval spots, and so on. When analyzing the PSR J0030+0415 data, Miller et al. (2019) found no evidence for different temperatures within the same hot spot or more than three heated regions on the stellar surface. They found that a model with three different nonoverlapping uniform oval spots is preferred over any models in which the spots overlap, describes the observed energy-resolved pulse profile adequately, and gives a fit to the data that is slightly, but not significantly, better than the best-fit model with two nonoverlapping, uniform oval spots. In the parameter estimation codes used by Miller et al. (2019) and Riley et al. (2019), the simpler spot models are subsets of the more elaborate spot models. An important aspect of the modeling approaches is that the complexity of the hot-spot arrangement is increased until a configuration that provides a satisfactory description of the data is reached. Miller et al. (2019) increased the complexity of the configuration still further to confirm that more complex spot patterns do not provide significantly better descriptions of the data. Despite the different approaches used by Miller et al. (2019) and Riley et al. (2019) in modeling the heated regions on the surface of PSR J0030+0451, the marginalized constraints on M and R found using these different approaches are consistent with one another, with 68% credible regions that have similar dimensions (spread/area) and largely overlap. This suggests that the mass and radius estimates obtained using these two approaches are relatively insensitive to the differences between them, which included different assumed priors on parameters (such as the stellar mass and radius) that are common to both models.

2.3. Relativistic Ray Tracing

To model the propagation of photons from the surface of a rapidly rotating MSP to Earth, we use the oblate Schwarzschild (OS) approximation proposed by Morsink et al. (2007) and developed further by AlGendy & Morsink (2014; see also Silva et al. 2021). In this approximation, the spacetime in the vicinity of the star is assumed to be the Schwarzschild spacetime, but the rotation-induced oblateness of the stellar surface is explicitly taken into account using analytical approximations to describe the shape of the surface and the effective surface gravity as a function of colatitude. This model also incorporates all of the relativistic effects that are important for the MSPs that are NICER targets for determining NS masses and radii, including gravitational light bending, redshift, and lensing, as well as propagation time delays, Doppler boosts, and aberration. Comparing the predictions of this approximation with the results of exact numerical computations shows that it is accurate to within $\sim 0.1\%$ for spin frequencies $\lesssim 300$ Hz (see Bogdanov et al. 2019b), which is more than adequate for our current parameter estimations.

The details of this model, its implementation, and the verification tests of the codes we have used to analyze the NICER pulse profile data are described in Paper II (Bogdanov et al. 2019b). In brief, given the mass M and equatorial radius R_{eq} of the NS and the locations, sizes, shapes, and temperature

distributions of the X-ray emitting regions on the stellar surface, the observed flux is computed as a function of photon energy and the rotational phase of the star. The resulting set of model pulse profiles is then compared with the energy-resolved pulse profile observed using NICER to estimate M and R_{eq} .

In the mass and radius parameter estimation analysis presented in Miller et al. (2019), the light-bending integral, lensing factor, and propagation time-delay integral are evaluated for a given dimensionless stellar radius (R/M) and photon emission angle θ and a particular combination of the star’s rotational phase ϕ and the colatitude θ_c of the surface element being considered by interpolating in a set of precomputed lookup tables. These tables cover R/M values in the range [3.1, 8.0], which includes the values of R/M for all expected M and R combinations, and θ values in the range [0, 2.2] rad. Only photon trajectories that do not intersect the oblate stellar surface are considered. In Riley et al. (2019), the deflection angle and propagation time delays at a set of colatitudes on every surface are computed on the fly, and spline interpolations of these are then used, which also yields the lensing factor. The shape of the stellar surface is parameterized by M , R_{eq} , and Ω . The arrays are bounded by the analytic most inward-going ray and the radial ray for every colatitude. Interpolation in precomputed lookup tables dramatically reduces the computational cost of producing the vast number of energy-resolved pulse profiles that are needed for the sampling algorithms used for Bayesian inference.

Multiple imaging—For stars that are sufficiently compact, there can be ray paths from a given emitting region to the observer that proceed directly to the observer or encircle the star one or more times, causing the observer to see multiple images of the emitting region. In order to be sure that our Bayesian analyses could find such compact configurations if any of the NICER MSPs are this compact, we cross-verified the ray paths and resulting pulse profiles for very compact stars that are given by the X-PSI²² package (referred to as the Amsterdam code in Paper II) and the Illinois–Maryland ray-tracing code.

The most compact, theoretically stable, nonrotating stars that are possible when the matter in them has a nonnegative isotropic pressure have $c^2R/GM = 9/4$ (the “Buchdahl limit”; see Buchdahl 1959). We first verified our ray-tracing codes by computing the photon number flux as a function of energy that would be observed from a nonrotating star with $c^2R/GM = 3.0$, emitting blackbody radiation uniformly from its entire surface. Our codes agreed with the analytical result for this case to within a few parts in 10^5 . We then computed the pulse profiles produced by two stellar models with rotation frequencies $\nu_{\text{rot}} = 346$ Hz, distance 1.0 kpc, $M = 2.14 M_{\odot}$, and $R_{\text{eq}} = 9.6$ km so that $c^2R_{\text{eq}}/GM = 3.038$. One model assumed two identical, antipodal, uniform-temperature circular hot spots, while the other assumed two nonidentical, nonantipodal, uniform-temperature circular hot spots. In the nonantipodal comparison that we feature in Figure 2, the observer inclination is $\theta_{\text{obs}} = 1.0$ rad. The first spot has a center at $\theta_{c1} = \pi/2$ rad, an angular radius $\Delta\theta_1 = 0.2$ rad, and an effective temperature measured at the surface of $T_{\text{eff},1} = 10^{6.1}$ K; the second spot has

²² Available for download at <https://github.com/ThomasEdwardRiley/xpsi>. For X-PSI, internal cross-checks of the “star-to-observer” ray-tracing approach were conducted by validating against a light-curve integrator that discretizes a moderately distant image plane (the so-called “observer-to-star” ray-tracing method).




Figure 2. Comparison of the pulse profiles at four different photon energies computed by the X-PSI (black) package and the Illinois–Maryland (red) ray-tracing code for a highly compact ($R_{\text{eq}}/M \approx 3.0$) star rotating at 346 Hz with two nonidentical, nonantipodal, uniform-temperature circular hot spots; see text for additional details. Each panel is labeled with the observed energy and the maximum number flux F_0 (in units of photons $\text{cm}^{-2} \text{s}^{-1} \text{keV}^{-1}$) over all phases at that energy.

a center at $\theta_{c2} = \pi - 1$ rad, an angular radius $\Delta\theta_2 = 0.1$ rad, and an effective temperature measured at the surface of $T_{\text{eff},2} = 10^{6.3}$ K. The longitude of the center of spot 2 is 0.55 cycles ahead of the longitude of the center of spot 1. The pulse profiles given by the two ray-tracing codes agree to within 2% or better in every pulse-phase bin and energy channel where the number flux is significant. Even at 2 keV, where the flux is more than 5 orders of magnitude smaller than the flux at 0.1 keV, the pulse profiles given by the two codes agree with one another to within 5% or better at all pulse phases. This difference is a small fraction of the statistical uncertainties in the pulse profiles that have been observed to date using NICER.

2.4. Interstellar Absorption

As X-ray photons propagate from the source to the observer, a fraction of them experience photoelectric absorption and scattering by the intervening ISM. The observed spectral flux $f(E)$ is described by the expression $f(E) = \exp^{-\sigma_{\text{ISM}} N_{\text{H}}} f_0(E)$, where $f_0(E)$ is the emitted spectral flux; σ_{ISM} is the attenuation per unit column density of hydrogen, which depends on the composition and ionization state of the ISM; and N_{H} is the column density of hydrogen along the line of sight from the source to the observer. To obtain the attenuation of the X-ray flux between our target MSPs and the NICER instrument for a given N_{H} , we use the state-of-the-art TBabs model for σ_{ISM} and elemental abundances in the ISM (Wilms et al. 2000).

In computing models of the observed energy-resolved pulse profiles, Miller et al. (2019) and Riley et al. (2019) used a precomputed lookup table that provides the attenuation at each NICER energy channel for a particular fiducial value of N_{H} . Because the attenuation factor at every energy scales the same

way with N_{H} , Miller et al. (2019) simply scaled the attenuation factor at every energy to obtain the value appropriate for the value of N_{H} being considered at each step in the fitting process.

The energy dependence of σ_{ISM} alters the observed X-ray spectrum of an MSP and therefore could, in principle, affect estimates of its mass and radius made using the NICER data. However, the MSPs that are targets for M and R determinations using NICER have exceptionally low values of N_{H} ($\sim 10^{20} \text{cm}^{-2}$), so interstellar extinction has only a small effect on their spectra. Furthermore, the pulse profile fitting method for estimating M and R uses the full energy-resolved profile and is therefore relatively insensitive to changes in the observed spectrum produced by interstellar extinction, which are pulse phase-independent. Thus, the constraints on the masses and radii of these MSPs obtained using this method are insensitive to uncertainties in the chemical composition of the ISM along the line of sight (see, e.g., Bogdanov et al. 2016), unlike constraints on the masses and radii of NSs in quiescent low-mass X-ray binaries that are based solely on their spectra and typically have to take into account high absorbing columns ($\gtrsim 10^{21} \text{cm}^{-2}$) along the line of sight (see Heinke et al. 2014).

2.5. Instrument Performance

In order to compare an energy-dependent model pulse profile with a pulse profile observed using NICER, the performance of the telescope must be taken into account. The NICER data consist of pulse-height events (not all of which are produced by X-ray photons from the source, or even by X-ray photons) in the N energy channels of the detector. The spectrum registered by the instrument, which is reported in units of counts per second per channel, is related to the photon spectrum $f(E)$ incident on the telescope, which is usually expressed in units of photons per second keV^{-1} , by the convolution equation

$$f_{\text{obs}}(N) = \int f(E) R(N, E) dE. \quad (1)$$

Here $R(N, E)$ is the instrument response expressed as the probability that an incoming photon of energy E will be detected in channel N . As is also the case for most other current X-ray instruments, the response of NICER’s X-ray Timing Instrument (XTI) is described by an RMF and an ARF. The RMF describes the probability that a photon with a given energy that is incident on the instrument is counted in each of its output channels. The ARF describes the efficiency of the instrument in counting incident photons of a given energy, expressed as the effective area (in square centimeters) of the instrument as a function of energy, and reflects the performance of the concentrator optics in focusing the incident X-rays, as well as the quantum efficiency of the detector.

In order to compare a model spectrum $f(E)$ with the spectrum $f_{\text{obs}}(N)$ registered by the XTI, the model spectrum must be convolved with the instrument response, as described by Equation (1). We have verified this procedure by comparing the results given by the independently developed pulse profile modeling codes described in Paper II to ensure that it is carried out correctly. We have also verified this procedure by comparing the results it gives for simple power-law and blackbody incident spectra with the results given by XSPEC (Arnaud 1996), when the same instrument response files are used.

2.6. Backgrounds

The final component needed to fit a pulse profile model to the NICER data is a model of the background, which we define as all counts accumulated by the NICER XTI that do not come from the X-ray emitting hot spots on the NS surface. The models that best fit the PSR J0030+0415 pulse profile data collected by NICER indicate that a significant fraction of the counts do not come from the hot spots (Miller et al. 2019; Riley et al. 2019), consistent with prior expectations of the source and background properties in XTI observations. These additional counts can be attributed to emission from gas in the immediate vicinity of the pulsar, sky background (from unrelated neighboring sources and the diffuse X-ray background), and noncosmic photons and particles in NICER’s local environment. Both Miller et al. (2019) and Riley et al. (2019) assumed that these nonspot background counts are not modulated at the rotation frequency of the pulsar, and the number of such counts in each energy bin was treated as a free parameter.²³ For future parameter estimation analyses of rotation-powered MSPs using NICER, we are developing the capability to incorporate information about the NICER nonsource background emission estimated using the methods described in (Paper I), as well as from any available archival X-ray observations of the same target with imaging telescopes such as XMM-Newton and Chandra.

2.7. Verification of Numerical Codes

In our numerical computations of model pulse profiles, we used the most up-to-date values and all significant figures of the relevant physical and astrophysical constants (e.g., Planck’s constant, the Boltzmann constant, the parsec, and GM_{\odot}) published by the Particle Data Group.²⁴

A critical requirement for producing accurate and reliable estimates of the masses and radii of the target MSPs is that the numerical codes used to compute model pulse profiles must be free of errors and have sufficient numerical precision. We sought to assure this by completing a set of code comparison and pulse profile verification tests using several independently developed codes to compute the model pulse profiles. As demonstrated in Paper II (Bogdanov et al. 2019b), we obtained excellent agreement between the results given by these different codes for the same assumed input parameters and between their results and a number of known analytical and semianalytical results.

We made sure that the spacing between the entries in the lookup tables we used to compute the spectra and beaming patterns of the radiation from our model atmospheres and to perform ray tracing in the vicinity of the NS models was fine enough that interpolation in these tables did not introduce numerical errors larger than the $\sim 0.1\%$ fractional error we established for the pulse profile comparisons in Paper II and for the model pulse profiles we use to perform our Bayesian parameter estimation.

²³ For details about how this background was modeled in Miller et al. (2019), see Section 3.4 of that paper. For details about how it was modeled in Riley et al. (2019), see Appendix B of Riley (2019).

²⁴ See <http://pdg.lbl.gov/2017/reviews/rpp2017-rev-phys-constants.pdf> and <http://pdg.lbl.gov/2017/reviews/rpp2017-rev-astrophysical-constants.pdf>.

2.8. Model Fitting and Parameter Estimation

By following the steps described above, we constructed energy-resolved models of the thermal X-ray pulse profiles produced by rotation-powered MSPs. We then used Bayesian techniques (see, e.g., Lo et al. 2013; Miller & Lamb 2015; Miller et al. 2019; and Riley et al. 2019 for details) to estimate the best-fit values and credible intervals of all of the parameters in a given pulse profile model, including the mass and radius of the pulsar, by fitting these models to the pulse profiles observed using NICER. In the next section, we describe some of the tests we performed to validate the algorithms and codes we used in this procedure.

3. Verification of Parameter Estimation Algorithms

Bayesian parameter estimation is widely used in astrophysics, and there are a number of publicly available software packages that can be used to sample the space defined by the model parameters and construct the posterior distribution for each parameter or any combination of parameters (e.g., Feroz et al. 2009; Foreman-Mackey et al. 2013; Handley et al. 2015; Brewer & Foreman-Mackey 2016; Betancourt 2017). The inference analysis performed by Riley et al. (2019) on the PSR J0030+0451 NICER data used the MultiNest nested sampling algorithm (Feroz et al. 2009) via the X-PSI package, while the analysis performed by Miller et al. (2019) used a hybrid approach in which MultiNest was used to perform an initial analysis, followed by more detailed sampling using PT-emcee, a parallel-tempered Markov Chain Monte Carlo algorithm (Foreman-Mackey et al. 2013). In early tests, the collaboration also considered using a basic Metropolis algorithm, an evolutionary algorithm (Stevens et al. 2016) coupled with a Metropolis algorithm, and PolyChord (Handley et al. 2015), but found them to be less well suited for the problem at hand.

Given the posterior distribution, the credible region (or regions) that contains a specified total probability can be defined uniquely in any number of dimensions as the minimum-volume region containing that probability. This is the approach used in the Illinois–Maryland code and the X-PSI²⁵ package to compute the two-dimensional credible regions reported in Miller et al. (2019) and Riley et al. (2019). In these codes, the minimum volume is computed by integrating over the posterior from higher to lower probability densities until the specified total probability is reached. To compute the one-dimensional posteriors reported in Miller et al. (2019) and Riley et al. (2019), both groups constructed the cumulative marginalized probability distribution for a given parameter and then reported probability intervals that are symmetric around the median. For example, the 90% credibility interval ranges from the 5th to the 95th percentile of the cumulative distribution.

As explained in Section 2.2, a realistic model of an MSP X-ray pulse profile produced by multiple hot spots necessarily requires a large number of free parameters, some of which are partially degenerate. Fitting such a model to NICER pulse profile data and extracting parameter estimates is therefore both algorithmically and computationally challenging. Moreover, the priors for the model parameters need to be chosen carefully,

²⁵ The X-PSI interfaces with the GetDist package (Lewis 2019) to calculate the minimum-area (or highest-density) credible regions in two dimensions.

so as not to produce specious posteriors.²⁶ While the publicly available sampling and inference software packages have been thoroughly verified and used widely, they possess advantages and disadvantages, depending on the particular problem in question (see, e.g., Table 1 of Brewer & Foreman-Mackey 2016, for a helpful summary). For instance, with an increasing number of parameters, certain algorithms may become very computationally inefficient, while for degenerate parameters and/or highly multimodal distributions, some algorithms may struggle with recovering all modes of the true posterior distribution. In light of this, it is crucial to verify whether the algorithms used for our parameter estimation with the NICER data are implemented correctly and capable of producing trustworthy inferences with the available tools and computational resources. To verify the ability of our algorithms and codes to overcome these problems, we conducted a number of tests. In them, our codes were used to analyze a series of progressively more complex synthetic NICER pulse profile data sets that were generated using models of rotating NSs with various numbers of hot spots. The codes we used for analysis of the NICER data of PSR J0030+0451 in Miller et al. (2019) and Riley et al. (2019) have been subjected to these tests and chosen based on their performance relative to other samplers.

In the first two subsections that follow, we describe in detail the results we obtained for two of these tests, which were carried out early in our development of parameter estimation procedures, prior to the analyses of the PSR J0030+0451 NICER data set. In these tests, our parameter estimation codes were used to analyze two different sets of synthetic NICER pulse profile data, one produced by a model pulsar with a single uniform-temperature circular hot spot and the other produced by a model pulsar with two different uniform-temperature circular hot spots. As we show below, the analysis procedures used by both Miller et al. (2019) and Riley et al. (2019)²⁷ to analyze these two synthetic pulse profiles produce estimates of the values of the model parameters and their credible regions that are consistent with the parameter values used to generate the synthetic data.

In the third subsection that follows, we describe two additional checks that Miller et al. (2019) performed to test the performance of their inference code by using it to analyze synthetic pulse profiles. The first of these additional checks was performed by analyzing a third synthetic pulse profile that was specially constructed to mimic the pulse profile of PSR J0030+0451 that was observed using NICER. This synthetic profile was generated using a model pulse profile produced by two nonoverlapping, oval hot spots and fit using a model with two possibly overlapping oval hot spots. As we discuss below, this fit gave estimated values and credible regions for all of the 15 parameters in the model that are consistent with the values that were used to generate the synthetic profile (see Section 4.2 of Miller et al. 2019).

The second additional check was made by performing posterior predictive assessments of whether the models with two and three oval hot spots used by Miller et al. (2019) are able to accurately describe the PSR J0030+0451 pulse profile and unmodulated background observed using NICER. As we detail below, both models pass this test.

We note that due to the complexity of the models we are using and the consequent high computational cost of analyzing the NICER data, at present it is not feasible to conduct a large-scale simulation-based calibration (see Berry et al. 2015; Talts et al. 2018; Riley 2019) of our inference procedures to demonstrate that our algorithms are capable of thoroughly exploring the vast parameter space and obtaining sufficient credible region coverage. The tests and cross-comparisons we have conducted do demonstrate that (1) the model components discussed in Section 2 have been correctly implemented in our data-fitting algorithms and (2) our inference codes yield parameter estimates and credible regions that are consistent with the values assumed in generating several different sets of synthetic pulse profile data.

3.1. An NS with a Single Uniform-temperature Circular Hot Spot

In the first test of our algorithms, we analyzed a synthetic NICER pulse profile produced by a spinning NS with a single uniform-temperature circular hot spot. In constructing this synthetic pulse profile, we used the following procedure.²⁸

1. We chose a gravitational mass $M = 1.4 M_{\odot}$ and a circumferential equatorial radius $R_{\text{eq}} = 12.0$ km (thus $c^2 R_{\text{eq}} / GM = 5.8047$).
2. We assumed a stellar rotational frequency $\nu_{\text{rot}} = 600$ Hz, as seen by a distant static observer. The shape of the stellar surface and the effective gravity at the stellar surface as a function of rotational colatitude were computed using the AlGendy & Morsink (2014) model.
3. We chose a distance to the NS of $D = 200$ pc.
4. We assumed that the observer is situated in the plane defined by the star's rotational equator (observer colatitude $\theta_{\text{obs}} = 90^\circ$).
5. We used the McPHAC nonmagnetic, fully ionized hydrogen atmosphere model (Haakonsen et al. 2012). As noted previously, while we have identified issues with the McPHAC code, for the purposes of this comparison of parameter estimation codes, they are unimportant, as both codes made use of the exact same atmosphere model lookup tables.
6. We considered a circular, uniform-temperature spot centered on the star's rotational equator (colatitude of the spot center $\theta_c = 90^\circ$) with an angular radius $\Delta\theta = 1.0$ rad and an effective temperature $kT_{\text{eff}} = 0.231139$ keV, as measured in the surface comoving frame.
7. For simplicity, in this simulation, we neglected absorption by the ISM (i.e., we set $N_{\text{H}} = 0$ cm⁻²).
8. We considered only the counts in NICER energy channels 0–299 inclusive (corresponding approximately to the energy band 0.1–3.09 keV).
9. For each of these 300 energy channels, we computed a high-accuracy profile of the pulse produced by the hot spot at 32 phase points (rather than integrating high-accuracy pulse profiles computed at closely spaced phases over each of the 32 phase bins).

²⁶ See Miller et al. (2019) and Riley et al. (2019) for a description of the priors used in the analyses of the NICER data of PSR J0030+0451.

²⁷ This test used a development version of the X-PSI package that preceded v0.1.

²⁸ The assumed parameter values for this test are largely the same as those assumed in test OS1b presented in Paper II. The synthetic data set used in this analysis, the tabulated atmosphere model, and the NICER instrument response files are available for download from Zenodo.

10. We folded these high-accuracy synthetic pulse profiles through the prelaunch NICER response matrix and then normalized this synthetic “observed” pulse profile so that the total expected number of counts in channels 0–299 (i.e., the expected number of counts from the hot spot) was 10^6 . This set the total observation time. This synthetic observed pulse profile was then sampled to produce a set of synthetic pulse profile data.
11. We assumed a phase-independent (unpulsed) background with a power-law photon number spectrum $dN/dE \propto E^{-2}$. Next, we folded this assumed energy-dependent but phase-independent background through the prelaunch NICER response matrix to produce the synthetic observed background. We then normalized this background so that the total expected number of background counts in channels 0–299 was 10^6 . (The expected total number of counts from the spot and the unpulsed background was therefore $10^6 + 10^6 = 2 \times 10^6$.) Finally, the normalized observed background was sampled to produce the synthetic background data.

We analyzed the synthetic pulse profile data (which included the synthetic background data) by fitting “practical” model pulse profiles²⁹ to them using the parameter estimation algorithms developed by Miller et al. (2019) and Riley et al. (2019). Both the synthetic pulse profile data and the model pulse profiles were binned in 32 phase bins,³⁰ which is a sufficient number of bins to accurately describe the smoothly varying pulse profile produced by the single-spot pulsar model considered in this test. The independent model profiles that were fit to the synthetic profile data using the X-PSI package and the Illinois–Maryland code had nine free parameters. Eight of these parameters occur in both models, namely, the star’s mass M and equatorial radius R_{eq} , the colatitude of the center of the hot spot θ_c , the angular radius $\Delta\theta$ and effective temperature T_{eff} of the hot spot, the angle θ_{obs} between the rotation axis of the star and the direction to the observer, the distance D to the observer, and the hydrogen column density N_{H} of the ISM along the line of sight to the observer. In addition, the X-PSI analysis included as a free parameter the azimuthal phase of the hot spot, whereas the Illinois–Maryland analysis marginalized over the overall phase of the pulse, making a total of nine individual free parameters in both analyses. Both model pulse profiles also included an arbitrary, phase-independent additive background parameter in each energy channel.

The resulting joint posterior distributions of R_{eq} and M given by the Illinois–Maryland code and the X-PSI package and the corresponding one-dimensional posterior probability distributions for R_{eq} and M are shown in Figure 3. This figure shows that (1) these two independent analyses of the same synthetic data set produced results that are statistically indistinguishable

²⁹ In principle, the codes used for these analyses can compute model pulse profiles to high precision using the OS approximation but at increased computational cost. Computing these profiles dominates the time required to evaluate the likelihood. Consequently, it is necessary to strike a balance between precision and speed. Thus, when analyzing these synthetic pulse profile data and the pulse profile data collected using NICER, we used so-called “practical” model profiles that have a fractional precision of only $\sim 0.1\%$ but can be generated in ~ 1 s per processor core (see Paper II for more details).

³⁰ From previous comparisons, we find that integrating over a phase bin or calculating the counts (after folding) per unit phase at the center of a phase bin and then multiplying by the width of the phase bin give answers that differ by an amount that is small compared with the uncertainty in the pulse profile data. Thus, either approach may be used, but it is helpful to specify which is assumed.

and (2) the credible regions given by both algorithms are consistent with the M and R_{eq} values that were assumed when generating the synthetic data, in that the injected values of M and R_{eq} belong to a posterior mode. Moreover, the residuals computed by comparing the best-fit pulse profile models to the pulse profile data do not show any statistically significant discrepancies.

3.2. An NS with Two Different Uniform-temperature Circular Hot Spots

As a second test of our analysis algorithms, we analyzed a synthetic NICER pulse profile produced by a spinning NS with two different uniform-temperature circular hot spots. The properties and locations of these two spots were chosen to produce a synthetic NICER pulse profile that qualitatively resembles the observed profile of PSR J0437–4715 (see, e.g., Figure 2 of Paper I). In constructing this synthetic pulse profile, we used the following procedure.

1. We chose a gravitational mass $M = 1.44 M_{\odot}$ and a circumferential equatorial radius $R_{\text{eq}} = 13.0$ km (thus $c^2 R_{\text{eq}}/GM = 6.1138$).
2. We assumed a stellar rotational frequency $\nu_{\text{rot}} = 173.6$ Hz, as seen by a distant static observer. The shape of the stellar surface and the effective gravity at the stellar surface as a function of colatitude were again computed using the AIGendy & Morsink (2014) model.
3. We chose a distance to the NS of $D = 156.3$ pc.
4. We assumed that the inclination of the observer’s sight line relative to the stellar spin axis is $\theta_{\text{obs}} = 0.733$ rad.
5. We used the NSX nonmagnetic, fully ionized hydrogen atmosphere model (Ho & Lai 2001; Ho & Heinke 2009).
6. We considered a circular, uniform-temperature, primary spot centered at colatitude $\theta_{c1} = 0.6283$ rad and azimuth $\phi = 0$ rad, where the zero of the azimuth is defined by the plane containing the observer, with an angular radius $\Delta\theta_1 = 0.01$ rad and an effective temperature $kT_{\text{eff},1} = 0.231139$ keV, as measured in the surface comoving frame.
7. We considered a circular, uniform-temperature, secondary spot centered at colatitude $\theta_{c2} = 2.077$ rad and azimuth $\phi_2 = 3.5343$ rad $= 0.5625$ cycles, with angular radius $\Delta\theta_2 = 0.33$ rad and an effective temperature $kT_{\text{eff},2} = 0.0577846$ keV, as measured in the surface comoving frame.
8. We included absorption by the ISM, assuming $N_{\text{H}} = 2 \times 10^{20}$ cm^{-2} and using the elemental abundances given by Wilms et al. (2000).
9. We assumed a total simulated observation time of 10^6 s.
10. We again assumed a pulse phase-independent background with the power-law spectrum $dN/dE \propto E^{-2}$, prior to folding through the NICER prelaunch instrument response matrix, normalized to a total of 10^6 expected counts in energy channels 25–299 inclusive after folding.

When generating the synthetic pulse profile data analyzed in this test and computing the model pulse profiles used to fit these synthetic data, we used the following ancillary information: (i) lookup tables for the spectra and beaming patterns produced by the NSX nonmagnetic, fully ionized hydrogen atmosphere model (Ho & Lai 2001; Ho & Heinke 2009), used to compute the properties of the radiation from the hot spots; (ii) the lookup table described in Section 2.4 for the TBabs

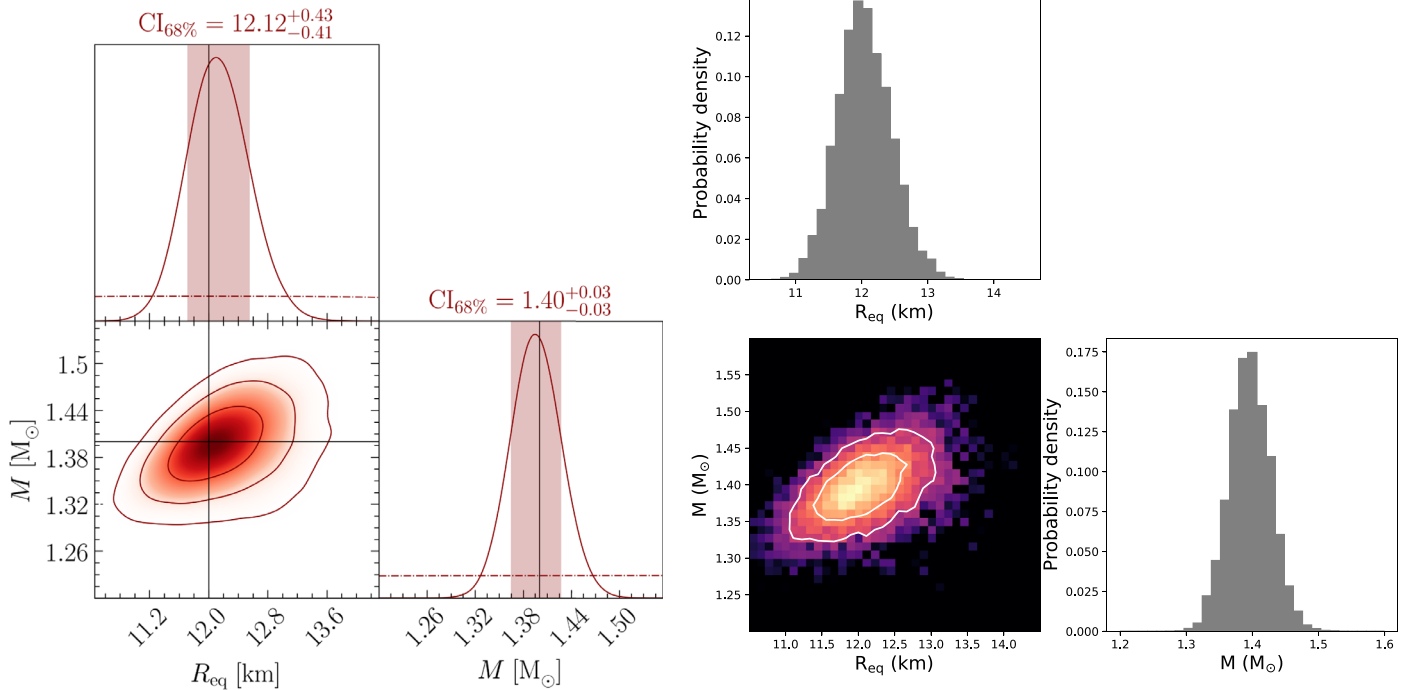


Figure 3. Posterior distributions in the M – R_{eq} plane obtained when the same synthetic NICER pulse profile produced by the stellar model with a single hot spot described in Section 3.1 is analyzed independently using the X-PSI package (left) and the Illinois–Maryland code (right). In the left panel, the red contours in the M – R_{eq} plot indicate the two-dimensional 1σ , 2σ , and 3σ credible regions, while the pink bands show the 1σ ranges of the one-dimensional posterior distributions. The black lines show the values of M and R_{eq} that were assumed in generating the synthetic data set. The dashed–dotted lines in the one-dimensional plots show the prior, and the vertical lines show the parameter values used to generate the synthetic data. In the right panel, the M – R_{eq} plot shows the joint posterior probability distribution, colored by credibility interval; the white contours enclose the 1σ and 2σ credible regions. The gray histograms show the one-dimensional posterior probability density distributions for R_{eq} and M . The 1σ credible region for R_{eq} is $12.03^{+0.42}_{-0.39}$, while the 1σ credible region for M is 1.40 ± 0.03 . These two independent analyses produce results that are statistically indistinguishable and consistent with the M and R_{eq} values assumed in generating the synthetic data. The residuals computed by comparing the best-fit pulse profile models with the synthetic pulse profile data are acceptable.

model, used to calculate the ISM absorption; and (iii) the publicly available, postlaunch (v1.02) NICER RMF and ARF, corrected for the off-axis pointing used when observing PSR J0437–4715,³¹ used to compute the pulse profile that would be observed using NICER. We again binned the synthetic pulse profile data and the model pulse profiles in 32 phase bins, which is a sufficient number of bins to accurately describe the smoothly varying rotation-induced flux modulation produced by the two-spot pulsar model considered in this test.

Due to the increased complexity of the two-spot pulse profile model relative to the single-spot model, we performed several preliminary checks on the two-spot model before proceeding to analyze the synthetic two-spot data. In these checks, we used the pulse profile modeling codes described in Paper II to independently produce and then cross-compare the six intermediate model pulse profiles we now list. The first four preliminary profiles were lists of the number of photons $\text{cm}^{-2} \text{s}^{-1}$ in each of the 32×700 pulse-phase–energy–channel bins that describe the profile; the final two profiles were lists of the number of counts expected in each NICER pulse-phase–energy–channel bin. These preliminary profiles were as follows.

1. The individual unsampled pulse profiles produced separately by the first and second hot spots, assuming no background and $N_{\text{H}} = 0$.
2. The unsampled combined pulse profile produced by both spots, again assuming no background and $N_{\text{H}} = 0$.
3. The unsampled combined pulse profile produced by both spots with ISM absorption included, assuming $N_{\text{H}} = 2.0 \times 10^{20} \text{ cm}^{-2}$.
4. The unsampled combined pulse profile produced by both spots with ISM absorption included, assuming $N_{\text{H}} = 2.0 \times 10^{20} \text{ cm}^{-2}$, and including the power-law background. The ISM absorption was not applied to the background.
5. The unsampled combined pulse profile produced by both spots after absorption by the ISM and the assumed background, after folding through the NICER off-axis detector response. The result is a list of the number of counts in each of the 32 phase bins for NICER energy channels 25–299 inclusive.
6. Finally, the realistic synthetic data set, consisting of the combined pulse profile produced by both spots after absorption by the ISM and the assumed background, after folding through the NICER off-axis detector response and sampling. The result is again a list of the number of counts in each of the 32 phase bins for NICER energy channels 25–299 inclusive, but now sampled. These are the data that were analyzed for this parameter estimation exercise.

³¹ As discussed in Paper I, due to the presence of a bright active galactic nucleus (AGN) near the position of PSR J0437–4715, a 1.47° pointing offset was used when observing this pulsar using NICER to reduce contamination of the pulsar data by photons from the AGN.

For the four preliminary pulse profiles that were not folded through the detector response matrix, the photon flux in each energy bin was determined by integrating over each bin, rather than computing the photon flux at the center of each bin. Each bin had a width of 0.005 keV; the low-energy edges of the bins were chosen to be 0.1, 0.105, 0.11, ..., 3.590, and 3.595 keV. The preliminary and final synthetic pulse profiles, expressed as phase-channel data, instrument response matrices, and tables of the extinction as a function of N_H are available on Zenodo.

We again analyzed the realistic synthetic pulse profile data by fitting “practical” model pulse profiles to these data using the parameter estimation codes developed by Miller et al. (2019) and Riley et al. (2019). The model profiles used by both groups have 13 free parameters. Twelve are present in both models: the star’s mass M and equatorial radius R_{eq} ; the colatitude of the center of the primary spot θ_{c1} ; the angular radius $\Delta\theta_1$ and effective temperature $T_{\text{eff},1}$ of the primary spot; the colatitude of the center of the secondary spot θ_{c2} ; the azimuthal phase ϕ_2 , angular radius $\Delta\theta_2$, and effective temperature $T_{\text{eff},2}$ of the secondary spot; the angle θ_{obs} between the rotation axis of the star and the direction to the observer; the distance D to the observer; and the hydrogen column density N_H of the ISM along the line of sight to the observer. In addition, the X-PSI analysis included the azimuthal phase ϕ_1 of the primary spot, whereas the Illinois–Maryland analysis marginalized over the overall phase of the pulse profile (an equivalent approach that makes the fitting procedure more efficient), making a total of 13 individual free parameters in both analyses. Both model pulse profiles also included an arbitrary, phase-independent background in each energy channel.

The resulting joint posterior distributions of R_{eq} and M given by the Illinois–Maryland code and the X-PSI package and the corresponding one-dimensional posterior probability distributions for R_{eq} and M are shown in Figure 4. This figure shows that (1) these two independent analyses produced results that are statistically indistinguishable and (2) the M and R_{eq} values that were assumed when generating the synthetic data belong to a posterior mode for both algorithms. Moreover, comparing the best-fit pulse profile models to the pulse profile data shows no statistically significant discrepancies in the residuals. This finding offers assurance that the posterior computations are well suited for the analyses of actual NICER data.

3.3. Analyses of Synthetic Pulse Profiles that Mimic the NICER Data on PSR J0030+0451

The analyses of the PSR J0030+0451 NICER data set by Miller et al. (2019) and Riley et al. (2019) revealed greater complexity in the hot-spot shapes and arrangement on the stellar surface compared to the early synthetic data analyses described above. This raises the question of whether the parameter estimation procedures of Miller et al. (2019) and Riley et al. (2019) give correct posterior probability distributions for the range of pulse profile models applied to the NICER data on PSR J0030+0451. To address this, in addition to testing their pulse profile modeling and inference procedures by analyzing the two synthetic pulse profiles produced by one and two circular spots as just described, Miller et al. (2019) also tested them by analyzing a third synthetic pulse profile that was specifically constructed to mimic the pulse profile of PSR J0030+0451 observed using NICER. We summarize this analysis and its results in the first paragraphs of this subsection.

We then describe posterior predictive assessments of the fits of the two pulse profile models used by Miller et al. (2019) to the NICER data on PSR J0030+0451. Unlike the previous analyses of synthetic pulse profiles, these assessments include both the PSR J0030+0451 pulse profile and the specific unmodulated thermal-like background observed using NICER.

Analyses of synthetic pulse profile data generated using a model with two oval hot spots—In addition to their analyses of the two synthetic pulse profiles produced by one and two circular spots just described, Miller et al. (2019) tested their inference procedure further by analyzing a third synthetic pulse profile, specifically constructed to mimic the pulse profile of PSR J0030+0451 observed using NICER. These synthetic data were produced using a pulse profile model with two different nonoverlapping uniform oval hot spots, with the model’s 14 parameters chosen to mimic the pulse profile of PSR J0030+0451 and assuming the NICER instrument’s response. These synthetic data were then fit using a 14-parameter model pulse profile that assumed two possibly different and overlapping uniform oval hot spots. The resulting fit gave estimated values and credible regions for all 14 of the parameters in the model that are consistent with the values of these parameters that were used to generate the synthetic profile (see also Section 4.2 in Miller et al. 2019).

Specifically, Miller et al. (2019) found that the values of 12 of the 14 parameters assumed when generating the synthetic pulse profile are within the $\pm 1\sigma$ credible intervals derived from their posterior probability distributions and that the assumed values of the other two parameters are within the corresponding $\pm 2\sigma$ credible intervals (see Table 5 in Miller et al. 2019). Thus, the one-dimensional posterior probability distributions that were obtained by analyzing the synthetic data are all consistent with the values of the model parameters that were assumed when the synthetic pulse profile data were produced. Miller et al. (2019) also found that the joint probability density distributions for M and R_{eq} obtained by fitting this model to the synthetic NICER pulse profile data are consistent with the values of the stellar mass and equatorial radius assumed when the synthetic data were produced (see Figure 3 in Miller et al. 2019).

Posterior predictive assessments of the fits of pulse profile models to the NICER data on PSR J0030+0451—In addition to the verification provided by the parameter estimation analyses of the three synthetic NICER pulse profile data sets that we discussed previously, Miller et al. (2019) provided information about the best fits of their models with two and three oval hot spots to the NICER data on PSR J0030+0451 that allow us to effectively perform posterior predictive assessments of the ability of their models, which had (a) the 205 Hz spin frequency of PSR J0030+0451, (b) two or three oval hot spots, and (c) an arbitrary unmodulated background, to accurately describe the PSR J0030+0451 pulse profile and the specific unmodulated thermal-like background observed using NICER.

A posterior predictive assessment of a fit to the data is the Bayesian version of a goodness-of-fit test (see, e.g., Gelman et al. 1996). It tests whether the log likelihood of the actual data, given the model and a set of assumed values of the parameters in the model, is or is not far out on the low tail of the distribution of log-likelihood values computed using a set of synthetic data; if it is, then according to this test, the model is not a good fit to the data.




Figure 4. Posterior distributions in the M – R_{eq} plane obtained when the same synthetic NICER pulse profile produced by the stellar model with two hot spots described in Section 3.2 is analyzed using the X-PSI package (left) and the Illinois–Maryland code (right). In the left panel, the black contours in the M – R_{eq} plot indicate the two-dimensional 1σ , 2σ , and 3σ credible intervals, while the gray bands show the 1σ ranges of the one-dimensional posterior distributions. The dashed lines in the one-dimensional plots show the priors, and the vertical lines show the parameter values used to generate the synthetic data. In the right panel, the M – R_{eq} plot shows the joint posterior probability distribution, colored by credibility interval; the white contours enclose the 1σ and 2σ credible regions. The gray histograms show the one-dimensional posterior probability density distributions for R and M . The 1σ credible region for R_{eq} is $13.45^{+0.66}_{-0.62}$ km, while the 1σ credible region for M is $1.44^{+0.06}_{-0.07}$ M_{\odot} . Again, these two independent analyses produce results that are statistically indistinguishable and consistent with the M and R_{eq} values assumed in generating the synthetic data. The residuals computed by comparing the best-fit pulse profile models with the synthetic pulse profile data are acceptable.

In applying this test to the fits of their models to the NICER data on PSR J0030+0451 obtained by Miller et al. (2019), we first note that the value of the phase-channel χ^2 reported by Miller et al. (2019) for the best fit of their pulse profile model with two oval spots to the PSR J0030+0451 data had a probability of 0.104 if this model was correct, while the best fit of their model with three oval spots had a probability of 0.120 if that model was correct. This means that the actual PSR J0030+0451 data collected by NICER could plausibly be drawn in a Poisson realization of their best-fitting model with two oval spots or from a Poisson realization of their best-fitting model with three oval spots and with the unmodulated backgrounds inferred from their fits.

Hence, for a specific set of the values of the parameters in either of their models, one can compute the log likelihood of the data given the model for that particular set of parameter values. Using the same model and the same set of parameter values, one can also generate synthetic data by Poisson sampling the expected number of counts in each phase-channel bin given by the model. One can then compute the log likelihood for the synthetic data given the model with those parameter values. Indeed, one can do this many times and thus obtain a distribution of the log likelihoods one would expect if the model were correct and the parameter values were the values assumed.

The information needed to perform this analysis was provided in Miller et al. (2019); one simply compares the $\pm 1\sigma$ and $\pm 2\sigma$ credible intervals for each of the parameters in their two pulse profile models with the best-fit values of these parameters, which are listed in their Tables 7 and 8. The

result is that 13 of the 15 parameters in the model with two oval hot spots fall in their respective $\pm 1\sigma$ intervals, while 15 of the 20 parameters in the model with three oval hot spots fall in their respective $\pm 1\sigma$ intervals. All of the parameters in the model with two oval hot spots fall in their respective $\pm 2\sigma$ intervals, and 19 of the 20 parameters in the model with three oval hot spots fall in their respective $\pm 2\sigma$ intervals. Thus, the best-fit models of Miller et al. (2019) easily pass this posterior predictive assessment, which assumes that the realizations of the synthetic data generated using the best-fit two- and three-spot models of Miller et al. (2019) exactly match the actual PSR J0030+0451 data collected by NICER and then analyzes that “synthetic” data using the models developed by Miller et al. (2019).

4. Dependence of M and R Constraints on Spin Rate and Hot-spot Size

As discussed in several previous works (see, e.g., Miller & Lamb 1998; Poutanen & Beloborodov 2006; Psaltis et al. 2014), the precision with which M and R_{eq} can be determined depends on many properties of the pulsar system. In particular, the colatitudes of hot spots and the angle between the sight line to the observer and the stellar spin axis strongly affect the precision. Geometries in which the spots and the sight line to the observer are near the plane of the star’s rotational equator provide tighter constraints than geometries in which either or both are close to the star’s spin axis (see, e.g., Bogdanov et al. 2008; Lo et al. 2013; Miller & Lamb 2015). In this section, we illustrate how the constraints on M and R_{eq} depend on the spin rate of the pulsar and the size of the hot spot by analyzing a set

of synthetic pulse profiles produced by pulsar models with a single circular, uniform-temperature hot spot and a phase-independent background but different spin rates and spot diameters.

To illustrate the strong dependence of the precisions of M and R_{eq} estimates on the spin rate and their much weaker dependence on the size of the hot spot, we generated synthetic NICER pulse profile data using pulsar models with the following four combinations of spin frequency ν_{rot} and hot-spot angular radius $\Delta\theta$, namely, case 1: $\nu_{\text{rot}} = 1$ Hz, $\Delta\theta = 0.01$ rad; case 2: $\nu_{\text{rot}} = 1$ Hz, $\Delta\theta = 1.0$ rad; case 3: $\nu_{\text{rot}} = 600$ Hz, $\Delta\theta = 0.01$ rad; and case 4: $\nu_{\text{rot}} = 600$ Hz, $\Delta\theta = 1.0$ rad. The values of the other parameters in the pulsar model were kept the same for all four cases. These parameters and their values were stellar equatorial radius $R_{\text{eq}} = 12$ km, stellar mass $M = 1.4 M_{\odot}$, angle between the sight line to the observer and the spin axis $\theta_{\text{obs}} = 1.047198$ rad (60°), colatitude of the center of the hot spot $\theta_c = 1.745329$ rad (100°), effective temperature of the hot spot $kT_{\text{eff}} = 0.231139$ keV, distance to the pulsar $D = 0.2$ kpc, and hydrogen column density $N_{\text{H}} = 1.0 \times 10^{20}$ cm $^{-2}$. For all four synthetic pulse profiles, we set the expected total number of counts from the hot spot to $\sim 360,000$ and the expected total number of background counts to $\sim 1,000,000$, approximately the number of counts from the hot spots and the phase-independent background obtained by Miller et al. (2019) in the best fit of their model with three oval spots to the NICER pulse profile data on PSR J0030+0451. We used a spectrum for the synthetic background that matches the spectrum of the phase-independent background found by Miller et al. (2019) in their best fit to the PSR J0030+0451 data.

Figure 5 shows the posterior distributions in the M – R_{eq} plane obtained by analyzing these four sets of synthetic data using the parameter estimation algorithm developed by Miller et al. (2019) and a pulse profile model with a single circular, uniform-temperature hot spot. These plots demonstrate how precisely the mass and radius of the star are encoded in the observed pulse profile by special and general relativistic effects.

If the star is rotating very slowly, Doppler boosts, aberration, and other special relativistic effects are unimportant, but gravitational lensing by the star—which depends only on M/R_{eq} —affects the pulse profile. Hence, fitting pulse profile models to the observed pulse profile of a slowly rotating star is expected to place a tighter constraint on M/R_{eq} than on M and R_{eq} separately. This is illustrated by case 1, which assumes a very slow (1 Hz) stellar rotation rate and a very small spot ($\Delta\theta = 0.01$ rad); in this case, the peak of the probability density roughly follows the expected $M/R_{\text{eq}} = 0.11 M_{\odot} \text{ km}^{-1}$ relation at low masses and radii but with a significant uncertainty in the slope at low masses and radii and a much larger uncertainty in the slope at high masses and radii. Case 2 also assumes a 1 Hz rotation rate but a much larger spot ($\Delta\theta = 1.0$ rad); in this case, the 1σ credible region is substantially larger than that for a very small spot and the probability density peak does not follow a line with $M/R_{\text{eq}} \approx \text{const}$. A spot this large is only gradually occulted by the body of the star as the star rotates, which probably explains why the slope of the probability density peak is steeper. In both cases, fitting the observed pulse profile constrains R_{eq} , but not very tightly.

If instead the star is rotating fairly rapidly ($\nu_{\text{rot}} \gtrsim 300$ Hz), the special relativistic effects produced by the velocity of the

star’s emitting surface relative to the observer—which is proportional to R_{eq} —affects the pulse profile. Hence, fitting pulse profile models to the observed pulse profile of a rapidly rotating star not only constrains M/R_{eq} but also constrains M and R_{eq} separately. This can be seen by comparing cases 1 and 2, which assume a very slow (1 Hz) rotation rate, with cases 3 and 4, which assume the same spot sizes as cases 1 and 2 but a much faster (600 Hz) rotation rate.³² As case 3 shows, for some rotation rates and spot sizes, the fractional uncertainty in the estimate of R_{eq} can be smaller than the fractional uncertainty in the estimate of M/R_{eq} (see also Psaltis et al. 2014; Miller & Lamb 2015; Stevens et al. 2016).

As noted above, the shape of the observed pulse profile also depends on the size of the hot spot. Radiation from a larger hot spot will be observed over a wider range of rotational phases, blurring the effects produced by the Doppler boost, aberration, and self-occultation of the hot spot by the star. We therefore expect that for a given number of counts from the spot and a given number of background counts, the credible region in the M – R_{eq} plane defined by a given cumulative probability will be somewhat larger, in general, for a larger hot spot. This is illustrated by case 1 compared to cases 2 and 3 compared to case 4. For all four synthetic data sets analyzed to produce Figure 5, the number of counts from the spot and the background were assumed to be the same, even though the larger spots have areas 10^4 times greater than the smaller spots. If all of these pulsars were observed for approximately the same length of time and all of the spots had the same temperature, $\sim 10^4$ times as many counts would be collected from the pulsars with the larger spots than from the pulsars with the smaller spots. If all else was equal, this would cause the constraints on M and R_{eq} to be $\sim 10^2$ times tighter for the stars with the larger spots than for the stars with the smaller spots.

5. Future Exploration of Possible Systematic Effects

Our simple description of hot spots cannot be entirely accurate. For example, the hotter regions on NSs will not be circles or ovals of perfectly uniform temperature. However, for the purpose of measuring the masses and radii of NSs observed using NICER, the models can be adequate without being accurate in every detail.

Indeed, Lo et al. (2013) and Miller & Lamb (2015) investigated the effects of systematic modeling errors on estimates of M and R_{eq} obtained by fitting X-ray pulse profile models to observations of the thermonuclear X-ray burst oscillations produced by some accreting MSPs or the X-ray pulsations produced by hot spots on the surfaces of nonaccreting MSPs. These investigations used pulse profile models to analyze synthetic pulse profiles in which the hot spot contributes a total of 10^6 counts and an unmodulated background contributes an additional 9×10^6 counts. This number of hot-spot counts is comparable to that which could be obtained by observing the brightest nonaccreting MSPs using NICER for about 10^6 s on source (corresponding roughly to exposure accumulated over a year of calendar time; see Paper I) or by combining observations of the oscillations of several dozen relatively bright X-ray bursts from the same accreting MSP (see Lo et al. 2013). This assumed ratio of oscillating

³² Using the OS approximation typically introduces errors of several percent for stellar spin frequencies of ~ 600 Hz, as explained in Paper II. However, errors of this size are unimportant for this purely illustrative comparison.




Figure 5. Posterior probability distributions in the M – R_{eq} plane obtained by analyzing synthetic NICER data from a pulsar with a single hot spot (see text) using the Illinois–Maryland code. The distributions are colored by credibility interval, as explained in Sections 3.1 and 3.2. The blue star marks the M and R_{eq} values assumed in generating the synthetic data; the white contours enclose the 1σ and 2σ credible regions. The four sets of synthetic pulse profile data that were analyzed were generated using the spin-frequency and spot-size combinations listed at the top of each panel. The panels are referred to in the text as case 1 (top left), case 2 (top right), case 3 (bottom left), and case 4 (bottom right). Note that the stellar compactness (M/R_{eq}) is not always the best-determined stellar property; indeed, for cases 3 and 4, the radius is determined with a better fractional precision than the compactness.

counts to background counts is conservative compared with the NICER data on PSR J0030+0415, for which the ratio of oscillating counts to background counts through 2019 January (1.8 million s of observing time) was about 360,000 to 1,000,000, rather than 1:9.

Importantly, these studies demonstrated that fitting rotating hot-spot models to energy-dependent X-ray pulse profiles gives results that are robust against several types of systematic effects. In particular, when pulse profile models were fit to synthetic data generated using spot shapes, energy spectra, temperature gradients, or surface beaming patterns different from those assumed in the models, the fits did not simultaneously produce (i) a statistically good fit (as indicated by the value of χ^2 obtained by comparing the best-fit bolometric pulse profile model with the bolometric pulse profile data), (ii) apparently tight constraints on the stellar mass and radius, and (iii) a significant bias in their inferred values. Thus, at least for these modeling errors, pulse profile analyses did not yield tight but misleading constraints.

There are other possible discrepancies between the model pulse profile used to analyze the observed X-ray pulse profile of a rotation-powered MSP and the actual profile of the pulsar. It would be worthwhile to further investigate their effects on

mass and radius estimates to determine whether or not they introduce significant biases. Possible discrepancies include the following.

Chemical composition of the atmosphere—The cosmic abundance of hydrogen and the expected rapid sinking of heavier elements in NS atmospheres justifies our standard assumption that the topmost layer of the NS atmosphere is pure hydrogen. In Miller et al. (2019), the NICER data from PSR J0030+0451 were also fit using a model with a nonmagnetic NSX helium atmosphere (Ho & Lai 2001). This analysis resulted in a best-fit mass of $2.7 M_{\odot}$, with most of the posterior probability above $2.0 M_{\odot}$. The tendency toward implausibly large NS masses suggests that an He atmosphere is unlikely, but this deserves further exploration, not only for PSR J0030+0451 but also for other MSPs that are NICER targets.

We also note that, since we are able to distinguish between helium and hydrogen atmosphere model results, which are caused by beaming patterns that differ at $\leq 5\%$ in the appropriate energy range (see Section 2.1), we expect to be sensitive to other atmosphere model effects that cause greater than 1%–2% differences in beaming patterns.

Magnetic fields—As discussed in Section 2.1, for the surface magnetic field strength inferred assuming a centered, purely dipolar stellar field for PSR J0030+0415, the effects of the magnetic field on the spectrum and beaming pattern of the radiation from the hot spots are negligible.

Kalopotharakos et al. (2021) investigated offset dipole plus quadrupole magnetic field configurations using static vacuum field and force-free global magnetosphere models aimed at reproducing the hot-spot patterns inferred for PSR J0030+0451 by Miller et al. (2019) and Riley et al. (2019). They demonstrated that by assuming the hot spots correspond to the open field line regions on the surface (i.e., the pulsar magnetic polar caps), the seemingly unusual spot geometries inferred for this pulsar from the NICER data can be reproduced very well by a number of plausible dipole plus quadrupole magnetic field arrangements. In addition, despite the presence of multipoles, the field is found to rarely exceed ~ 10 times the dipole spin field and is thus typically $\lesssim B_0$. Therefore, in the case of PSR J0030+0451, magnetic influences on radiation transport in the NS atmosphere are likely small in the NICER band. While strong magnetic fields are likely to produce beaming patterns that are incompatible with the energy-dependent pulse profile from PSR J0030+0451, the effect of different atmosphere beaming patterns on the NS mass and radius inferences deserves further investigation.

Depth of energy deposition—The surface thermal emission from rotation-powered MSPs is most likely caused by heating by backward-flowing energetic particles produced by pair cascades at the polar caps or in the outer magnetosphere and by global return currents³³ (e.g., Spitkovsky 2006; Harding & Muslimov 2011). In our models, we assume that the energy of the returning particles that bombard the NS surface is deposited many optical depths below the photosphere. To assess the validity of this assumption, note that plasma instabilities are unlikely to play a significant role in the stopping process (see, e.g., Timokhin & Arons 2013), that a returning electron or positron with a Lorentz factor $\gamma > 10^3$ has a total stopping depth³⁴ in hydrogen larger than 90 g cm^{-2} , and that it will therefore deposit the bulk of its energy at a density around $\rho \sim 10^3 \text{ g cm}^{-3}$ (Bethe & Heitler 1934; Tsai 1974). Given that the column depth of a typical nonmagnetic NS H atmosphere is $0.1\text{--}1 \text{ g cm}^{-2}$ (see, e.g., Zavlin et al. 1996), a particle with a Lorentz factor $\gamma > 10^3$ will deposit negligible energy in such an atmosphere. Lower-energy particles, with $\gamma = 10\text{--}10^3$, would stop at depths of $2\text{--}90 \text{ g cm}^{-2}$ in H and therefore could potentially deposit some of their energy at densities as low as 10 g cm^{-3} . In strong magnetic fields, the stopping depth of electrons is reduced by magneto-Coulomb interactions in which the electrons are excited to higher Landau levels in the Coulomb fields of ions and then de-excite by emitting synchrotron or cyclotron radiation (Ho et al. 2007; Potekhin 2014). However, the MSPs that are NICER targets are thought to have surface magnetic fields too weak for this process to significantly decrease the stopping depths of electrons or positrons.

Bauböck et al. (2019) investigated the radiative properties of an NS atmosphere (assumed to be gray) heated by returning

particles that deposit their energy in the upper layers of the atmosphere. They found that this occurs for low-energy particles with $\gamma \sim 2\text{--}10$. Salmi et al. (2020) performed a similar analysis but treated Compton scattering within the atmosphere exactly. They concluded that if most of the returning particles have Lorentz factors $\gamma \lesssim 100$, their energy will be deposited at relatively low optical depths, causing the atmosphere to develop a hotter, optically thin skin on top of a cooler, optically thick region.

Timokhin & Arons (2013) found that most of the heating of the atmosphere by returning particles is caused by those with $\gamma \gg 10\text{--}100$, and in some cases as high as $\sim 10^7$ (see especially Figures 17 and 18 in Timokhin & Arons 2013). For MSPs in particular, the lowest-energy particles in calculated pair spectra (see Harding & Muslimov 2011, Figure 10) have $\gamma \sim 10^3\text{--}10^4$. This is because the very low magnetic fields of MSPs require much higher photon energies to create pairs. Even though, for the time-dependent cascades, there are some low-energy thermal particles moving downward, their number densities are $100\text{--}1000$ times lower than the higher-energy pairs, so they are unlikely to contribute much of the heating. It therefore seems probable that our standard assumption that the energy of returning particles is deposited at large optical depths is correct for the MSPs observed with NICER. Nevertheless, it is important to examine the effect that small deviations in the atmosphere model beaming pattern might have on the inferred values of the NS parameters.

Nonspot background emission—For the NICER MSP data under consideration for $M - R$ parameter estimation analyses, a substantial fraction of the observed emission does not originate from the hot regions on the stellar surface and is not modulated at the pulsar spin period. As described in Paper I, there are a number of sources of background radiation that can affect NICER X-ray data, which include particle and photon background from the local low-Earth-orbit environment of NICER, the unresolved diffuse X-ray background, other point sources that fall within the telescope field of view, and other emission processes in the vicinity of the pulsar under consideration. In the analyses of the NICER data of PSR J0030+0451 presented in Miller et al. (2019) and Riley et al. (2019), no estimated background is explicitly taken into account; this is the most conservative approach. Instead, non-hot-spot emission in each detector channel is treated as a free parameter and assumed to have no dependence on spin phase. Thus, all background emission that is not modulated at the pulsar period is incorporated in a straightforward fashion.

It is important to emphasize that although such an assumption introduces a large number of additional free parameters, this would not result in bias, as the best solution is still reachable by the model. Instead, it would lead to greater imprecision in the measurement, i.e., broader credible intervals for the inferred parameters compared to a model that can accurately describe the observed pulse profile but contains a smaller number of parameters. This is demonstrated by the analysis in Miller et al. (2019) of the synthetic NICER data that mimicked that of PSR J0030+0451 (including the estimated background level), which also assumed no knowledge of the background emission and produced estimated values and credible regions consistent with the values assumed in constructing the synthetic pulse profile.

By the same token, the inclusion of independent estimates of the nonspot contribution to the observed X-ray emission as a

³³ These pulsars have characteristic ages much greater than the age of $\sim 10^6$ yr when the emission from the surface of an NS produced by its initially hot interior could exceed the emission produced by backward-flowing particles and return currents.

³⁴ <https://physics.nist.gov/PhysRefData/Star/Text/ESTAR.html>

function of energy in the parameter estimation analyses can potentially lead to significantly tighter constraints on the NS mass and radius estimates, provided that the background estimates are reliable and do not introduce measurement bias by over/underestimating the true background. As our understanding of instrumental and other sources of background improves, we anticipate exploiting such knowledge as informative priors in our Bayesian inferences, with potentially significant implications for the precision and accuracy of the resulting NS mass and radius estimates. As noted in Section 2.6, we are developing the capability to incorporate complementary spectral information about the non-source emission that accounts for a substantial fraction of the events collected in NICER observations of the MSP targets under consideration. This includes using the background estimated from the “3c50” and space weather–based techniques developed by the NICER team (see Paper I), as well as existing X-ray imaging data from other observatories. The details of the implementation of background information will be described in subsequent publications.

Pulsed nonthermal emission—While all backgrounds that are not modulated at the rotation frequency of the pulsar are included in the analysis algorithms being used by Miller et al. (2019) and Riley et al. (2019), any modulated nonthermal emission from the pulsar in the energy band that is currently being analyzed (0.3–3 keV) is not included in our models. For PSR J0437–4715, there is marginal evidence for pulsations in the 2–20 keV band (Guillot et al. 2016), but the photon flux from this apparently nonthermal emission is only $\sim 3\%$ of the total flux at energies < 2 keV, where the modulated thermal component is prominent. Also, PSR J0030+0451 exhibits a power-law spectral tail evident above ~ 2 keV in XMM-Newton data (Bogdanov & Grindlay 2009) but is too faint to test for the presence of pulsations. It is therefore unlikely that this apparently nonthermal pulsed flux significantly affected our analysis of the pulsed thermal emission from PSR J0030+0451, but additional studies aimed at quantifying its impact on the NS mass and radius estimates are needed.

NICER instrumental calibration—Because the energy-dependent sensitivities of the NICER detectors are not perfectly known,³⁵ uncertainties in their responses could, in principle, affect estimates of the masses and radii of MSPs made using NICER data. Energy-independent uncertainties in the effective area were already incorporated in Miller et al. (2019) by adding extra effective uncertainty to the distance to the MSPs being observed and in Riley et al. (2019) by inclusion of calibration parameters (estimated via a procedure similar to that described by Ludlam et al. 2018 using observations of the Crab Nebula as a reference). The effect of energy-dependent uncertainties in the response of the NICER XTI could be more subtle. However, using different methods (see the individual papers for details), Miller et al. (2019) and Riley et al. (2019) found no evidence that systematic errors in the instrumental calibration biased their estimates of the mass and radius of PSR J0030+0451. The NICER calibration is being refined with time, so the evolution of mass–radius estimates as updated calibrations are used will provide information about whether calibration errors could significantly affect these estimates.

6. Conclusions

We have described the physical components of the model of pulsed thermal X-ray emission by MSPs that we have used and will continue to use to analyze NICER data on rotation-powered MSPs to infer their masses and radii. We also reported a comparison of model pulse profiles produced by a star with $R_{\text{eq}}/M \approx 3$ performed to cross-check the accuracy of these codes for future analyses of NICER pulse profile data, where considering more compact stars might be relevant.

We have also described our verification of the parameter estimation algorithms and codes used by Miller et al. (2019) and Riley et al. (2019) to analyze the NICER data on PSR J0030+0451 by using them to analyze synthetic pulse profile data produced by two pulsar models, one with a single hot spot and one with two hot spots. We showed that these codes yield credible regions for all of the parameters in the pulse profile models used to analyze the synthetic data that are consistent with each other and with the parameter values used to generate the synthetic pulse profile data. The NICER synthetic data sets and auxiliary files used in these comparisons are provided to the community on Zenodo ([10.5281/zenodo.4710158](https://zenodo.org/record/4710158)) as supplementary material to this article to allow testing of other independently developed codes.

We also provided an illustration of how the stellar spin rate and the size of its hot spot affect the constraints on the mass and radius of NSs that can be obtained using the pulse profile fitting method. The results obtained by using this technique to analyze synthetic NICER pulse profiles generated using four model pulsars highlight the importance of rapid spin for obtaining constraints on M and R_{eq} separately. In the case of hot-spot size, for equivalent source and background counts, NSs with smaller spots would, in principle, produce tighter M and R_{eq} constraints but in practice would require substantially longer exposures compared to stars with larger spots to achieve this.

We discussed the assumptions made in our analyses concerning the composition of the NS atmosphere, the strength and geometry of its surface magnetic field, and the depth at which returning particles deposit their energy. We concluded that these assumptions are reasonable and based on our current best understanding of the surface properties of rotation-powered MSPs as informed by multiwavelength observations and theory. Nevertheless, further study of the aforementioned effects may be useful to ascertain their impact on the NS mass and radius estimates. Finally, we argued that the estimates of the mass and radius of PSR J0030+0451 that have been made using NICER are not significantly affected by the current approach to modeling the unpulsed background, the possible pulsation of its nonthermal medium-energy X-ray emission, and that instrumental calibration errors are unlikely to introduce significant biases in our mass and radius estimates. We concluded that these assumptions are likely accurate, though more exhaustive studies of some may be useful.

The results presented here provide additional confidence that the current measurements of the mass and radius of PSR J0030+0451 and future measurements of the masses and radii of this MSP and others made using NICER will be reliable as well as precise.

This work was supported by NASA through the NICER mission and the Astrophysics Explorers Program. S.B. was funded in part by NASA grants NNX17AC28G and

³⁵ Due to the challenge of obtaining suitable reference targets for precise calibration measurements, all observations with X-ray telescopes suffer from an uncertainty in the absolute effective area across all energies, which is commonly estimated to be at the level of $\pm 10\%$.

80NSSC20K0275. W.C.G.H. appreciates the use of computer facilities at the Kavli Institute for Particle Astrophysics and Cosmology and acknowledges support through grant 80NSSC20K0278 from NASA. A.J.D., M.C.M., and F.K.L. acknowledge the University of Maryland supercomputing resources that were made available for conducting the research reported in this letter. M.C.M. thanks the Radboud Excellence Initiative for supporting his stay at Radboud University and was also supported by a Visiting Researcher position at the Perimeter Institute for Theoretical Physics in the early stages of this project. T.E.R., D.C., and A.L.W. acknowledge support from ERC Starting grant No. 639217 CSINEUTRONSTAR and ERC Consolidator grant No. 865768 AEONS. This work was sponsored by NWO Exact and Natural Sciences for the use of supercomputer facilities and was carried out on the Dutch national e-infrastructure with the support of SURF Cooperative. Z.W. acknowledges support from the NASA postdoctoral program. S.G. acknowledges the support of the CNES. S.M. M. acknowledges support from NSERC. T.E.R., D.C., and A.L.W. are grateful to W. M. Farr and I. Mandel for introducing us to the notion of simulation-based calibration of posterior computation and for their advice and suggestions about weaker cross-checks, all of which informed our perspective on what level of calibration we attained. 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 the ArXiv.

Facility: NICER.

Software: Python/C language (Oliphant 2007), GNU Scientific Library (Gough 2009), Cython (Behnel et al. 2011), SciPy (Virtanen et al. 2020), NumPy (van der Walt et al. 2011), Matplotlib (Hunter 2007), MultiNest (Feroz et al. 2009), emcee (Foreman-Mackey et al. 2013), GetDist (Lewis 2019).

ORCID iDs

Slavko Bogdanov  <https://orcid.org/0000-0002-9870-2742>
 Alexander J. Dittmann  <https://orcid.org/0000-0001-6157-6722>
 Wynn C. G. Ho  <https://orcid.org/0000-0002-6089-6836>
 Frederick K. Lamb  <https://orcid.org/0000-0002-3862-7402>
 Simin Mahmoodifar  <https://orcid.org/0000-0003-2386-1359>
 M. Coleman Miller  <https://orcid.org/0000-0002-2666-728X>
 Sharon M. Morsink  <https://orcid.org/0000-0003-4357-0575>
 Thomas E. Riley  <https://orcid.org/0000-0001-9313-0493>
 Tod E. Strohmayer  <https://orcid.org/0000-0001-7681-5845>
 Anna L. Watts  <https://orcid.org/0000-0002-1009-2354>
 Devarshi Choudhury  <https://orcid.org/0000-0002-2651-5286>
 Sebastien Guillot  <https://orcid.org/0000-0002-6449-106X>
 Alice K. Harding  <https://orcid.org/0000-0001-6119-859X>
 Paul S. Ray  <https://orcid.org/0000-0002-5297-5278>
 Zorawar Wadiasingh  <https://orcid.org/0000-0002-9249-0515>
 Michael T. Wolff  <https://orcid.org/0000-0002-4013-5650>
 Craig B. Markwardt  <https://orcid.org/0000-0001-9803-3879>
 Keith C. Gendreau  <https://orcid.org/0000-0001-7115-2819>

References

- Alcock, C., & Illarionov, A. 1980, *ApJ*, 235, 534
 AlGendy, M., & Morsink, S. M. 2014, *ApJ*, 791, 78
 Arnaud, K. A. 1996, in ASP Conf. Ser. 101, *Astronomical Data Analysis Software and Systems V*, ed. G. H. Jacoby & J. Barnes (San Francisco, CA: ASP), 17
 Badnell, N. R., Bautista, M. A., Butler, K., et al. 2005, *MNRAS*, 360, 458
 Bauböck, M., Psaltis, D., & Özel, F. 2019, *ApJ*, 872, 162
 Baym, G., Hatsuda, T., Kojo, T., et al. 2018, *RPPH*, 81, 056902
 Behnel, S., Bradshaw, R., Citro, C., et al. 2011, *CSE*, 13, 31
 Beloborodov, A. M. 2002, *ApJL*, 566, L85
 Berry, C. P. L., Mandel, I., Middleton, H., et al. 2015, *ApJ*, 804, 114
 Betancourt, M. 2017, arXiv:1701.02434
 Bethe, H., & Heitler, W. 1934, *RSPSA*, 146, 83
 Bildsten, L., Salpeter, E. E., & Wasserman, I. 1992, *ApJ*, 384, 143
 Bilous, A. V., Watts, A. L., Harding, A. K., et al. 2019, *ApJL*, 887, L23
 Blaes, O. M., Blandford, R. D., Madau, P., & Yan, L. 1992, *ApJ*, 399, 634
 Bogdanov, S., & Grindlay, J. E. 2009, *ApJ*, 703, 1557
 Bogdanov, S., Grindlay, J. E., & Rybicki, G. B. 2008, *ApJ*, 689, 407
 Bogdanov, S., Guillot, S., Ray, P., et al. 2019a, *ApJL*, 887, L25
 Bogdanov, S., Heinke, C. O., Özel, F., & Güver, T. 2016, *ApJ*, 831, 184
 Bogdanov, S., Lamb, F. K., Mahmoodifar, S., et al. 2019b, *ApJL*, 887, L26
 Braje, T. M., Romani, R. W., & Rauch, K. P. 2000, *ApJ*, 531, 447
 Brewer, B. J., & Foreman-Mackey, D. 2016, arXiv:1606.03757
 Brown, E. F., Bildsten, L., & Chang, P. 2002, *ApJ*, 574, 920
 Buchdahl, H. A. 1959, *PhRv*, 116, 1027
 Cadeau, C., Morsink, S. M., Leahy, D., & Campbell, S. S. 2007, *ApJ*, 654, 458
 Chang, P., & Bildsten, L. 2003, *ApJ*, 585, 464
 Chang, P., & Bildsten, L. 2004, *ApJ*, 605, 830
 Chen, A. Y., Yuan, Y., & Vasilopoulos, G. 2020, *ApJL*, 893, L38
 Colgan, J., Kilcrease, D. P., Magee, N. H., et al. 2016, *ApJ*, 817, 116
 Contopoulos, I., Kazanas, D., & Fendt, C. 1999, *ApJ*, 511, 351
 Contopoulos, I., & Spitkovsky, A. 2006, *ApJ*, 643, 1139
 Feroz, F., Hobson, M. P., & Bridges, M. 2009, *MNRAS*, 398, 1601
 Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, 125, 306
 Gänsicke, B. T., Braje, T. M., & Romani, R. W. 2002, *A&A*, 386, 1001
 Gelman, A., Meng, X.-L., & Stern, H. 1996, *Statistica Sinica*, 6, 733
 Gendreau, K. C., Arzoumanian, Z., Adkins, P. W., et al. 2016, *Proc. SPIE*, 9905, 99051H
 González-Caniulef, D., Guillot, S., & Reisenegger, A. 2019, *MNRAS*, 490, 5848
 Gough, B. 2009, *GNU Scientific Library Reference Manual* (Devon: Network Theory Ltd.)
 Gralla, S. E., Lupsasca, A., & Philippov, A. 2017, *ApJ*, 851, 137
 Guillot, S., Kaspi, V. M., Archibald, R. F., et al. 2016, *MNRAS*, 463, 2612
 Haakonsen, C. B., Turner, M. L., Tacik, N. A., & Rutledge, R. E. 2012, *ApJ*, 749, 52
 Hameury, J. M., Heyvaerts, J., & Bonazzola, S. 1983, *A&A*, 121, 259
 Handley, W. J., Hobson, M. P., & Lasenby, A. N. 2015, *MNRAS*, 450, L61
 Harding, A. K., & Muslimov, A. G. 2011, *ApJ*, 743, 181
 Hebel, K., Lattimer, J. M., Pethick, C. J., & Schwenk, A. 2013, *ApJ*, 773, 11
 Heinke, C. O. 2013, *JPhCS*, 432, 012001
 Heinke, C. O., Cohn, H. N., Luggner, P. M., et al. 2014, *MNRAS*, 444, 443
 Heinke, C. O., Rybicki, G. B., Narayan, R., & Grindlay, J. E. 2006, *ApJ*, 644, 1090
 Ho, W. C. G., & Heinke, C. O. 2009, *Natur*, 462, 71
 Ho, W. C. G., Kaplan, D. L., Chang, P., van Adelsberg, M., & Potekhin, A. Y. 2007, *MNRAS*, 375, 821
 Ho, W. C. G., & Lai, D. 2001, *MNRAS*, 327, 1081
 Hunter, J. D. 2007, *CSE*, 9, 90
 Iglesias, C. A., & Rogers, F. J. 1996, *ApJ*, 464, 943
 Kalapotharakos, C., Harding, A. K., Kazanas, D., & Wadiasingh, Z. 2019, *ApJL*, 883, L4
 Kalapotharakos, C., Wadiasingh, Z., Harding, A. K., & Kazanas, D. 2021, *ApJ*, 907, 63
 Lai, D. 2001, *RvMP*, 73, 629
 Lattimer, J. M., & Prakash, M. 2001, *ApJ*, 550, 426
 Lattimer, J. M., & Prakash, M. 2005, *PhRvL*, 94, 111101
 Lattimer, J. M., & Prakash, M. 2016, *PhR*, 621, 127
 Lewis, A. 2019, arXiv:1910.13970
 Lo, K. H., Miller, M. C., Bhattacharyya, S., & Lamb, F. K. 2013, *ApJ*, 776, 19
 Lockhart, W., Gralla, S. E., Özel, F., & Psaltis, D. 2019, *MNRAS*, 490, 1774
 Ludlam, R. M., Miller, J. M., Arzoumanian, Z., et al. 2018, *ApJL*, 858, L5
 McClintock, J. E., Narayan, R., & Rybicki, G. B. 2004, *ApJ*, 615, 402

- Mészáros, P. 1992, *Theoretical Astrophysics* (Chicago: Univ. Chicago Press)
- Mihalas, D. 1978, *Stellar Atmospheres* (2nd ed.; Sn Francisco, CA: W. H. Freeman)
- Miller, M. C. 2013, in *Timing Neutron Stars: Pulsations, Oscillations and Explosions*, ed. T. Belloni, M. Méndez, & C. Zhang (Berlin: Springer), 1
- Miller, M. C., & Lamb, F. K. 1998, *ApJL*, **499**, L37
- 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
- Morsink, S. M., Leahy, D. A., Cadeau, C., & Braga, J. 2007, *ApJ*, **663**, 1244
- Nättilä, J., & Pihajoki, P. 2018, *A&A*, **615**, A50
- Oertel, M., Hempel, M., Klähn, T., & Typel, S. 2017, *RvMP*, **89**, 015007
- Oliphant, T. E. 2007, *CSE*, **9**, 10
- Özel, F. 2013, *RPPPh*, **76**, 016901
- Özel, F., & Psaltis, D. 2009, *PhRvD*, **80**, 103003
- Page, D. 1995, *ApJ*, **442**, 273
- Pavlov, G. G., & Zavlin, V. E. 1997, *ApJL*, **490**, L91
- Pechenick, K. R., Ftaclas, C., & Cohen, J. M. 1983, *ApJ*, **274**, 846
- Potekhin, A. Y. 2010, *A&A*, **518**, A24
- Potekhin, A. Y. 2014, *PhyU*, **57**, 735
- Poutanen, J., & Beloborodov, A. M. 2006, *MNRAS*, **373**, 836
- Poutanen, J., & Gierliński, M. 2003, *MNRAS*, **343**, 1301
- Psaltis, D., & Özel, F. 2014, *ApJ*, **792**, 87
- Psaltis, D., Özel, F., & Chakrabarty, D. 2014, *ApJ*, **787**, 136
- Raaijmakers, G., Greif, S. K., Riley, T. E., et al. 2020, *ApJL*, **893**, L21
- Raaijmakers, G., Riley, T. E., Watts, A. L., et al. 2019, *ApJL*, **887**, L22
- Rajagopal, M., & Romani, R. W. 1996, *ApJ*, **461**, 327
- Randhawa, J. S., Meisel, Z., Giuliani, S. A., et al. 2019, *ApJ*, **887**, 100
- Rauch, T., Suleimanov, V., & Werner, K. 2008, *A&A*, **490**, 1127
- Read, J. S., Lackey, B. D., Owen, B. J., & Friedman, J. L. 2009, *PhRvD*, **79**, 124032
- Riley, T. E. 2019, PhD thesis, Univ. Amsterdam <https://hdl.handle.net/11245.1/aa86fcf3-2437-4bc2-810e-cf9f30a98f7a>
- Riley, T. E., Watts, A. L., Bogdanov, S., et al. 2019, *ApJL*, **887**, L21
- Romani, R. W. 1987, *ApJ*, **313**, 718
- Rosen, L. C. 1968, *Ap&SS*, **1**, 372
- Salmi, T., Suleimanov, V. F., Nättilä, J., & Poutanen, J. 2020, *A&A*, **641**, A15
- Salmi, T., Suleimanov, V. F., & Poutanen, J. 2019, *A&A*, **627**, A39
- Silva, H. O., Pappas, G., Yunes, N., & Yagi, K. 2021, *PhRvD*, **103**, 063038
- Spitkovsky, A. 2006, *ApJL*, **648**, L51
- Steiner, A. W., Lattimer, J. M., & Brown, E. F. 2013, *ApJL*, **765**, L5
- Stevens, A. L., Fiege, J. D., Leahy, D. A., & Morsink, S. M. 2016, *ApJ*, **833**, 244
- Strohmayer, T. E. 1992, *ApJ*, **388**, 138
- Suleimanov, V., & Werner, K. 2007, *A&A*, **466**, 661
- Suleimanov, V. F., Klochkov, D., Pavlov, G. G., & Werner, K. 2014, *ApJS*, **210**, 13
- Suleimanov, V. F., Pavlov, G. G., & Werner, K. 2010, *ApJ*, **714**, 630
- Suleimanov, V. F., Pavlov, G. G., & Werner, K. 2012, *ApJ*, **751**, 15
- Talts, S., Betancourt, M., Simpson, D., Vehtari, A., & Gelman, A. 2018, [arXiv:1804.06788](https://arxiv.org/abs/1804.06788)
- Timokhin, A. N., & Arons, J. 2013, *MNRAS*, **429**, 20
- Tsai, Y.-S. 1974, *RvMP*, **46**, 815
- van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, *CSE*, **13**, 22
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *NatMe*, **17**, 261
- Watts, A. L., Andersson, N., Chakrabarty, D., et al. 2016, *RvMP*, **88**, 021001
- Weinberg, N., Miller, M. C., & Lamb, D. Q. 2001, *ApJ*, **546**, 1098
- Wijngaarden, M. J. P., Ho, W. C. G., Chang, P., et al. 2019, *MNRAS*, **484**, 974
- Wijngaarden, M. J. P., Ho, W. C. G., Chang, P., et al. 2020, *MNRAS*, **493**, 4936
- Wilms, J., Allen, A., & McCray, R. 2000, *ApJ*, **542**, 914
- Zavlin, V. E., Pavlov, G. G., & Shibano, Y. A. 1996, *A&A*, **315**, 141