



The Radius of PSR J0740+6620 from NICER and XMM-Newton Data

M. C. Miller¹, F. K. Lamb^{2,3}, A. J. Dittmann¹, S. Bogdanov⁴, Z. Arzoumanian⁵, K. C. Gendreau⁵, S. Guillot^{6,7}, W. C. G. Ho⁸, J. M. Lattimer⁹, M. Loewenstein^{5,10,11}, S. M. Morsink¹², P. S. Ray¹³, M. T. Wolff¹³, C. L. Baker¹⁴, T. Cazeau⁵, S. Manthripragada¹⁵, C. B. Markwardt¹⁶, T. Okajima⁵, S. Pollard⁵, I. Cognard^{17,18}, H. T. Cromartie^{19,29}, E. Fonseca^{20,21,22,23}, L. Guillemot^{17,18}, M. Kerr²⁴, A. Parthasarathy²⁵, T. T. Pennucci^{26,27}, S. Ransom²⁶, and I. Stairs²⁸

¹ Department of Astronomy and Joint Space-Science Institute, University of Maryland, College Park, MD 20742-2421 USA; miller@astro.umd.edu

² 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

³ Department of Astronomy, University of Illinois at Urbana-Champaign, 1002 West Green Street, Urbana, IL 61801-3074, USA

⁴ Columbia Astrophysics Laboratory, Columbia University, 550 West 120th Street, New York, NY 10027, USA

⁵ X-Ray Astrophysics Laboratory, NASA Goddard Space Flight Center, Code 662, Greenbelt, MD 20771, USA

⁶ IRAP, CNRS, 9 avenue du Colonel Roche, BP 44346, F-31028 Toulouse Cedex 4, France

⁷ Université de Toulouse, CNES, UPS-OMP, F-31028 Toulouse, France

⁸ Department of Physics and Astronomy, Haverford College, 370 Lancaster Avenue, Haverford, PA 19041, USA

⁹ Department of Physics and Astronomy, Stony Brook University, Stony Brook, NY 11794-3800, USA

¹⁰ Department of Astronomy, University of Maryland, College Park, MD 20742-2421, USA

¹¹ Center for Research and Exploration in Space Science and Technology, NASA/GSFC, Greenbelt, MD 20771, USA

¹² Department of Physics, University of Alberta, 4-183 CCIS, Edmonton, AB T6G 2E1, Canada

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

¹⁴ Applied Engineering and Technology Directorate, NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA

¹⁵ Instrument Systems and Technology Division, NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA

¹⁶ Astrophysics Science Division, NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA

¹⁷ Station de Radioastronomie de Nançay, Observatoire de Paris, CNRS/INSU, Université d'Orléans, F-18330, Nançay, France

¹⁸ Laboratoire de Physique et Chimie de l'Environnement, CNRS, 3A Avenue de la Recherche Scientifique, F-45071, Orléans Cedex 2, France

¹⁹ Cornell Center for Astrophysics and Planetary Science and Department of Astronomy, Cornell University, Ithaca, NY 14853, USA

²⁰ Department of Physics, McGill University, 3600 rue University, Montréal, QC H3A 2T8, Canada

²¹ McGill Space Institute, McGill University, 3550 rue University, Montréal, QC H3A 2A7, Canada

²² Department of Physics and Astronomy, West Virginia University, P.O. Box 6315, Morgantown, WV 26506, USA

²³ Center for Gravitational Waves and Cosmology, West Virginia University, Chestnut Ridge Research Building, Morgantown, WV 26505, USA

²⁴ Space Science Division, Naval Research Laboratory, Washington, DC 20375 5352, USA

²⁵ Max-Planck-Institut für Radioastronomie, Auf dem Hügel 69, D-53121 Bonn, Germany

²⁶ National Radio Astronomy Observatory, 520 Edgemont Road, Charlottesville, VA 22903, USA

²⁷ Institute of Physics, Eötvös Loránd University, Pázmány P.s. 1/A, 1117 Budapest, Hungary

²⁸ Department of Physics and Astronomy, University of British Columbia, 6224 Agricultural Road, Vancouver, BC V6T 1Z1, Canada

Received 2021 March 13; revised 2021 June 1; accepted 2021 June 4; published 2021 September 8

Abstract

PSR J0740+6620 has a gravitational mass of $2.08 \pm 0.07 M_{\odot}$, which is the highest reliably determined mass of any neutron star. As a result, a measurement of its radius will provide unique insight into the properties of neutron star core matter at high densities. Here we report a radius measurement based on fits of rotating hot spot patterns to Neutron Star Interior Composition Explorer (NICER) and X-ray Multi-Mirror (XMM-Newton) X-ray observations. We find that the equatorial circumferential radius of PSR J0740+6620 is $13.7^{+2.6}_{-1.5}$ km (68%). We apply our measurement, combined with the previous NICER mass and radius measurement of PSR J0030+0451, the masses of two other $\sim 2 M_{\odot}$ pulsars, and the tidal deformability constraints from two gravitational wave events, to three different frameworks for equation-of-state modeling, and find consistent results at ~ 1.5 – 5 times nuclear saturation density. For a given framework, when all measurements are included, the radius of a $1.4 M_{\odot}$ neutron star is known to $\pm 4\%$ (68% credibility) and the radius of a $2.08 M_{\odot}$ neutron star is known to $\pm 5\%$. The full radius range that spans the $\pm 1\sigma$ credible intervals of all the radius estimates in the three frameworks is 12.45 ± 0.65 km for a $1.4 M_{\odot}$ neutron star and 12.35 ± 0.75 km for a $2.08 M_{\odot}$ neutron star.

Unified Astronomy Thesaurus concepts: X-ray sources (1822); Millisecond pulsars (1062); Neutron stars (1108); Neutron star cores (1107)

1. Introduction

Neutron stars are unique laboratories for the study of dense matter. Their cores consist of matter that is believed to be catalyzed to the ground state, at a few times nuclear saturation density (a mass density $\rho_s \approx 2.7$ – 2.8×10^{14} g cm⁻³, or a baryonic number density $n_s \approx 0.16$ fm⁻³). The combination of high density and the expected large neutron–proton

asymmetry in neutron star cores cannot be duplicated in laboratories. Hence, observations of neutron stars can provide us with a valuable window into an otherwise inaccessible realm of nuclear physics.

Over the last several years great strides have been made in neutron star observations, and thus in our understanding of the equation of state (EOS: pressure as a function of energy density) of neutron star matter at high densities (see, e.g., Pavlov & Zavlin 1997; Bhattacharyya et al. 2005; Steiner et al. 2010; Miller 2013; Miller & Lamb 2016; Özel et al. 2016;

²⁹ NASA Hubble Fellowship Program Einstein Postdoctoral Fellow.

Nättilä et al. 2017 for earlier perspectives). Three neutron stars have a gravitational mass established to be $M \sim 2 M_{\odot}$: PSR 1614–2230 with $M = 1.908 \pm 0.016 M_{\odot}$ (the uncertainties here and below are for the 68% credible region) (Demorest et al. 2010; Fonseca et al. 2016; Arzoumanian et al. 2018); PSR J0348+0432 with $M = 2.01 \pm 0.04 M_{\odot}$ (Antoniadis et al. 2013); and PSR J0740+6620 with $M = 2.08 \pm 0.07 M_{\odot}$ (Cromartie et al. 2020; Fonseca et al. 2021). The existence of such high-mass neutron stars indicates that the EOS of neutron star matter is relatively hard, i.e., that it yields high pressures at a few times ρ_s . The lack of a clear signature of tidal deformation in gravitational wave observations of GW170817 (Abbott et al. 2017, 2018; De et al. 2018) and GW190425 (Abbott et al. 2020a) indicates that the EOS is not too hard. When taken together, these high measured masses and the upper limits on the tidal deformability derived from these gravitational wave observations have already narrowed significantly the range of allowed EOS models. The neutron star radius and mass measurements made using data from the Neutron Star Interior Composition Explorer (NICER) are potentially even more informative.

The first simultaneous measurements of the mass and radius of a neutron star made using NICER data were those of the millisecond pulsar PSR J0030+0451, which was determined to have a gravitational mass of $M \approx 1.44 \pm 0.15 M_{\odot}$ and an equatorial circumferential radius of $R_e \approx 13 \pm 1$ km (Miller et al. 2019; Riley et al. 2019; compare with the 11.9 ± 1.4 km radius inferred for the two $\sim 1.4 M_{\odot}$ neutron stars in the gravitational wave event GW170817; Abbott et al. 2018; De et al. 2018). Assuming that systematic errors in the NICER measurements are unimportant, and hence that the fractional uncertainty of each of these measurements decreases as the inverse square root of the observing time devoted to each pulsar, we expect these uncertainties to become $\sim 30\%$ – 40% smaller within the next few years. A precision this high will improve significantly our understanding of neutron star matter. It is also important to determine whether the masses and radii, and the EOS of neutron star matter, determined using NICER observations of other neutron stars are consistent with those determined using NICER observations of PSR J0030+0451.

In particular, it is valuable to measure the radii of higher-mass neutron stars. Such stars have higher central densities than a $\sim 1.4 M_{\odot}$ star, which means that a radius measurement will probe the EOS in a higher-density regime. For example, whereas the measurement of the radius of a $1.4 M_{\odot}$ star has its greatest impact on our understanding of matter at 1.6 times nuclear saturation density, the measurement of the radius of a $2.0 M_{\odot}$ star tells us primarily about matter at 2.2 times nuclear saturation density (see Section IV.D of Drischler et al. 2021; see also Xie & Li 2020 for additional perspectives on the importance of radius measurements for high-mass neutron stars). As pointed out in the context of the $\sim 1.9 M_{\odot}$ pulsar PSR J1614–2230, even a low-precision radius measurement of a high-mass pulsar, or indeed even a solid lower limit to the radius, will be useful in the construction of more accurate EOS for dense matter (Miller 2016).

Here we report our analysis of X-ray data for the $\sim 2.1 M_{\odot}$ pulsar PSR J0740+6620, which has a rotational frequency of 346.53 Hz (Cromartie et al. 2020). In contrast to PSR J0030+0451, which is an isolated neutron star, PSR J0740+6620 is in a binary and therefore we have independent measurements of the mass and orbital inclination. This improves the precision of

the fits. The challenge presented by PSR J0740+6620 is that its NICER count rate is only $\sim 5\%$ that of PSR J0030+0451. As a result, counts in NICER data from PSR J0740+6620 are dominated by sources unassociated with the pulsar, to a far greater degree than is true for PSR J0030+0451. This significantly reduces the modulation fraction and harmonic structure in the X-ray pulse waveform, which in turn reduces the information we have about the emitting regions (which we henceforth call “spots”) and the mass and radius of the star. If we had precise and reliable information about the background in NICER observations we could use this to improve our mass and radius estimates, but at present we need to turn to other sources of information.

In this Letter we therefore present an analysis of PSR J0740+6620 based on both NICER and X-ray Multi-Mirror (XMM-Newton) data. The NICER data provide information about the modulated emission from the star, which is produced from hot spots on the star as the star rotates. The XMM-Newton data have a far smaller rate of background counts than the NICER data and are therefore informative about the total flux from the star. Combined, the NICER and XMM-Newton data sets thus yield an improved measurement of the modulation fraction and the harmonic structure from PSR J0740+6620, and hence improve the precision with which we can measure the radius. In particular, stronger gravitational lensing allows more of the star to be visible, which lowers the modulation of the signal. Since gravitational lensing depends on the ratio of mass to radius, for fixed mass, an increase in the pulsed fraction means that a larger radius will be inferred if all other parameters are fixed.

In Section 2 we describe the selection of the NICER and XMM-Newton data, including blank-sky XMM-Newton data that have been accumulated over the duration of its mission and that help us refine further our estimate of the stellar flux. In Section 3 we discuss our analysis methods. These are based on the approaches described in previous NICER papers (Bogdanov et al. 2019, 2021; Miller et al. 2019), but are augmented here by the extra data sets and our inclusion of background information from the XMM-Newton blank-sky observations. In Section 4 we present our analysis results including tests of the adequacy of our models; we find that the radius of PSR J0740+6620 is $R_e = 13.7^{+2.6}_{-1.5}$ km at 68% credibility. We also describe several differences between our analysis and that in the parallel paper Riley et al. (2021), in Section 4.6. In Section 5 we present the implications of our PSR J0740+6620 analysis, combined with our previous results on PSR J0030+0451 and other data, for the EOS of cold dense matter. Our conclusions are in Section 6. We present our full corner plots and tables of posteriors in the appendices. Samples from our full posterior probability distributions are available at <https://zenodo.org/record/4670689>.

2. Observations

2.1. Data Used

We base our analysis of PSR J0740+6620 on NICER event data that included individual event times, event energies, and event pulsation phase. Our observing strategy and data reduction resulted in a data set with very low noise properties, as was necessary because of the faintness of this pulsar.

Our analysis of the X-ray data on PSR J0740+6620 comes from a series of 215 observations obtained with NICER XTI collected between 2018 September 21 and 2020 April 17, with

a net exposure time of 1602.68 ks (after all filtering was applied, as described below). The final event data were obtained using the following filtering criteria: (i) we only included ObsIDs when all 52 NICER detectors were active; (ii) we excluded events with energies outside the range 0.3–3.0 keV; (iii) we excluded events from DetID 34 (which frequently exhibits elevated count rates compared to the average); (iv) we only considered time intervals when PSR J0740+6620 was situated at least 80° from the Sun to minimize optical loading and scattered solar X-ray background photons; and (v) we rejected data obtained at low cut-off rigidities ($COR_SAX < 5$) to minimize particle background contamination. The filtered events were assigned pulse phases using the PSR J0740+6620 ephemeris from Fonseca et al. (2021). This event list ultimately consisted of 7334 separate good time intervals (GTIs). We then searched this event list for the pulsar signal. If the events in a particular GTI did not increase the pulsar signal to noise we rejected that GTI. The pulsar detection significance was highest in the energy range 0.31 to 1.18 keV. Ultimately, after all filtering and other exclusions were finished, we ended up with 521,004 NICER XTI events. The details of the NICER observations and pulse detection for PSR J0740+6620 can be found in Wolff et al. (2021).³⁰

We use the latest available calibration products, namely the nixtref20170601v002 redistribution matrix file and the nixtia-veonaxis20170601v004 ancillary response file (ARF). For the latter, the effective areas per energy channel were corrected by a factor of 51/52 to account for the removal of all events from DetID 34. The response properties of NICER incorporated in these calibration products represented the best effort by the instrument modeling team that were available in August 2020 when the analysis presented here was begun.

PSR J0740+6620 was also targeted with XMM-Newton on three separate occasions: 2019 October 26 (ObsID 0851181601), 2019 October 28 (ObsID 0851181401), and 2019 November 1 (ObsID 0851181501) as part of a Director’s Discretionary Time program. All three European Photon Imaging Camera (EPIC) instruments (pn, MOS1, and MOS2) were operated in “Full Frame” imaging mode with the “Thin” filters in the optical path. Due to the long detector read-out times (73.4 ms for EPIC-pn and 2.6 s for EPIC-MOS1/2), the data do not permit a pulse-phase-resolved analysis. Data summed over phase are used in the analysis that follows. The XMM-Newton data were processed with the Science Analysis Software (SAS³¹) using the recommended procedures. We employed the recommended event grade PATTERN and FLAG values that ensured only real photon events recorded by the instruments were included in our analysis (≤ 12 for MOS1/2 and ≤ 4 for pn). All three of these XMM-Newton observations were obtained while the satellite was in a part of its orbit that suffered from intense particle flaring effects. Only 6.8 ks of the total 20.8 ks of EPIC-pn exposure had sufficiently low background to be used in our analysis, and only 18.0/18.7 ks of the total 26.7/34.3 ks of MOS1/2 exposures could be utilized. A circular extraction region of radius $25''$ centered on the radio timing position of PSR J0740+6620 was used to produce clean source event lists. The XMM-Newton

instrument response files specific to the PSR J0740+6620 observations were then generated using the *rmfgen* and *arfgn* tools in SAS.

Under normal observing circumstances, one uses regions near the target source in the instrument image plane from the same observations to estimate the observational background. In our case, however, due to the short durations of these observations, very few background counts were available to estimate the background. As an alternative, we obtained representative background estimates with longer exposures from the blank-sky event files provided by the XMM-Newton Science Operations Centre. The blank-sky images were filtered in the same manner as the PSR J0740+6620 images and the background was extracted from the same location on the detector image plane as the pulsar. The resulting background spectrum was then rescaled so that the exposure times and pixel-area factors matched those of the PSR J0740+6620 exposures.

2.2. Calibration of the Instruments

The effective area of the NICER XTI is determined primarily using observations of the Crab pulsar and nebula. The energy-dependent residuals in the fits to the Crab spectrum are typically at the level of $\lesssim 2\%$.³² The calibration accuracies for the XMM-Newton EPIC-MOS and EPIC-pn cameras are determined to better than 3% and 2% (at 1σ), respectively.³³ However, in the absence of a suitable absolute calibration source, the uncertainties in the absolute energy-independent effective areas that we use for the NICER XTI and the three XMM-Newton detectors are estimated to be no more than $\pm 10\%$ (e.g., Figure C2 in Ricci et al. 2021; see Section 3.7 for further discussion).

3. Methods

Our approach to modeling the X-ray data is largely the same as it was in Miller et al. (2019), except that for our analysis of PSR J0740+6620 we used both NICER and XMM-Newton data, and we used blank-sky observations to estimate the non-source counts in the XMM-Newton data. In Section 3.1 through Section 3.3, we summarize briefly the approaches used in Miller et al. (2019), and refer the reader to that paper for details and tests of our procedures. In Section 3.4 we discuss the new aspects of our fitting procedure, which include analyzing multiple data sets and incorporating independent estimates of the background made using the XMM-Newton data. We conclude this section with short discussions of our parameter estimation procedure in Section 3.5 (which is the same as was described in Miller et al. 2019), our choice of energy channels in Section 3.6 (which is different from the choice made in Miller et al. 2019 because we are using an updated energy calibration for NICER and because we also use XMM-Newton data in our analysis), and the effects of a possible difference in the absolute calibration of NICER and XMM-Newton in Section 3.7.

³⁰ See https://heasarc.gsfc.nasa.gov/docs/nicer/mission_guide/ for the standard definitions of the observational terms used in this section.

³¹ The XMM-Newton SAS is developed and maintained by the Science Operations Centre at the European Space Astronomy Centre and the Survey Science Centre at the University of Leicester.

³² See <https://heasarc.gsfc.nasa.gov/docs/heasarc/caldb/nicer/docs/xti/NICER-xti20200722-Release-Notesb.pdf> for further details.

³³ See in particular Table 1 in <https://xmmweb.esac.esa.int/docs/documents/CAL-TN-0018.pdf>.

3.1. Modeling the Emission from the Stellar Surface

The soft X-ray emission from old, non-accreting pulsars such as PSR J0740+6620 is believed to be generated when particles with large ($\gg 100$) Lorentz factors are produced in pair cascades and deposit their energy deep in the atmosphere of the neutron star (Tsai 1974; Harding & Muslimov 2001, 2002; Bogdanov et al. 2007; see Bauböck et al. 2019 and Salmi et al. 2020 for explorations of the consequences for the beaming pattern of emission from the stellar surface if significant energy is in particles with lower Lorentz factors). The specific intensity as a function of the angle between the photon propagation direction and the local surface normal depends on the assumed composition, the ionization state, and the strength (and potentially the orientation) of the surface magnetic field at the point of emission.

We follow Miller et al. (2019) in assuming that the atmosphere is pure hydrogen and that the magnetic field can be neglected. We now discuss this in more detail.

Composition of the upper atmosphere. The surface gravity of a neutron star is great enough that the lightest element present is expected to float to the top within seconds to minutes (based on extrapolations of the calculations in Alcock & Illarionov 1980). As hydrogen is also the most abundant element in the universe and PSR J0740+6620 likely underwent prolonged accretion to reach its current rotational frequency, it is probable that the atmosphere is pure hydrogen (Blaes et al. 1992; Wijngaarden et al. 2019, 2020). Note that there are some accreting neutron star binaries in which the transferred matter is thought to have little to no hydrogen (such as 4U 1820–30; see, e.g., Strohmayer & Brown 2002), but such binaries have extremely short orbital periods (e.g., Stella et al. 1987 find that the orbital period of 4U 1820–30 is just 685 s). Thus the binary is very compact and as a result the hydrogen envelope is believed to have been completely stripped off. In contrast, the binary containing PSR J0740+6620 has an orbital period of 4.77 days (Cromartie et al. 2020) and the companion could therefore have had a significant hydrogen envelope. Even if the accreted elements were primarily heavier than hydrogen, spallation might create enough hydrogen to dominate the atmosphere (Bildsten et al. 1992; Randhawa et al. 2019).

Ionization fraction. The hot spots on PSR J0030+0451 were inferred to have effective temperatures (i.e., the temperature of a blackbody that would produce a bolometric photon flux equal to the bolometric photon flux produced by the model atmosphere we are considering, which is not a blackbody) $kT_{\text{eff}} > 0.1$ keV (Miller et al. 2019; Riley et al. 2019), which implies that essentially all atmospheric hydrogen is completely ionized. The fits we find for PSR J0740+6620 have effective temperatures $kT_{\text{eff}} \gtrsim 0.08$ keV for both spots, which also implies nearly full ionization. However, the atmospheric densities are high enough that some neutral atoms may be present. We therefore performed fits using the NSX code for model atmospheres assuming full ionization (Ho & Lai 2001) and also with partial ionization (Ho & Heinke 2009; see Badnell et al. 2005 for the relevant opacity tables). We find no significant difference in the inferred mass and radius between the results obtained using our two different ionization assumptions. This is consistent with previous work (Lo et al. 2013; Miller & Lamb 2015), which finds that in analyses of phase-channel X-ray data such as is collected using NICER, even if the assumptions used in the fit deviate from the properties of the star there is no significant bias in the mass and

radius if the fit appears to be statistically good (as measured by, e.g., χ^2/dof or other standard statistics). Here we report only on runs that allow for the possibility of partial ionization.

Effect of the magnetic field on the soft X-ray emission from the stellar surface. Using Equation (12) of Contopoulos & Spitkovsky (2006), assuming a magnetic field inclination of $\pi/2$ and $\alpha = 0$ for the magnetic field configuration, the period $P = 0.00289$ s and period derivative $\dot{P} = 1.219 \times 10^{-20}$ of PSR J0740+6620 (Cromartie et al. 2020) imply that for a centered dipole the polar magnetic field at the surface is $B \approx 3 \times 10^8$ G. The electron cyclotron energy at that field strength is ≈ 3 eV, which has a small effect on atomic structure (Lai 2001; Potekhin 2014) and is ~ 100 times lower than the lowest energies we consider in our fits. The spot configuration resulting from our analysis is close enough to a centered dipole that it is plausible that the effects of the field on radiative transfer are negligible. For example, in our best fit the two spots are separated by ≈ 2 radians on the surface. The chord length is therefore $\approx \sin(1) \approx 0.84$ times smaller than a diameter, so if the magnetic field is a dipole centered between the two spots then the field strength at the surface is $\approx 1/0.84^3 \approx 1.7$ times larger than it would be if the field were centered at the center of the star. Thus in this example the surface magnetic field at the spots would increase only to $\approx 5 \times 10^8$ G. We follow Miller et al. (2019) in noting that if the field is strong enough to change significantly the spectrum and beaming pattern, then fits assuming a negligible field would likely be statistically poor and/or would yield unreasonable parameter values. Therefore the good quality of our fits could be considered an argument in favor of our assumption that the field can be neglected. However, this conclusion is not certain and a rigorous resolution of this question would require the construction, and use in fits, of atmospheric models with a fine grid of magnetic field strengths and orientations to the local surface normal. Such models are challenging to compute at the required $B \sim 10^9$ – 10^{10} G because at those field strengths Coulomb and magnetic effects are comparable to each other.

3.2. Our Approach to Modeling the Soft X-Ray Waveform

Our modeling of the NICER pulse waveform of PSR J0740+6620 builds on the modeling of the waveform of PSR J0030+0451 discussed in Miller et al. (2019); see Bogdanov et al. (2019, 2021) for detailed descriptions of how we computed pulse waveforms and tested their accuracy. The waveform of PSR J0740+6620 has two clear peaks and therefore cannot be modeled using a star with a single circular hot spot. Unlike for PSR J0030+0451, and possibly related to the faintness of PSR J0740+6620 and hence the much smaller number of counts in its pulse waveform, we find that a model with two uniform circular hot spots produces a waveform that adequately fits the data, and that adding complexity to the temperature pattern on the stellar surface does not improve the fit. For example, adding a third hot spot to the model or allowing the spots to be oval rather than circular did not improve the fit. We therefore fit the NICER and XMM-Newton data for PSR J0740+6620 using a pulse waveform produced by two uniform-temperature circular spots.

Once the number of spots to be considered and their allowed shapes (e.g., oval or circular) are specified, our spot modeling algorithm automatically explores the full parameter space, to find the configurations with the highest likelihoods (see Miller et al. 2019). The algorithm not only considers configurations in

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

Parameter	Definition	Assumed Prior
$c^2 R_e / (GM)$	Inverse of stellar compactness	3.2–8.0
M	Gravitational mass	$\exp[-(M - 2.08 M_\odot)^2 / 2(0.09 M_\odot)^2]$
θ_{c1}	Colatitude of spot 1 center	0 to π radians
$\Delta\theta_1$	Spot 1 radius	0–3 radians
$kT_{\text{eff},1}$	Spot 1 effective temperature	0.011–0.5 keV
$\Delta\phi_2$	Spot 2 longitude difference	0–1 cycles
θ_{c2}	Colatitude of spot 2 center	0 to π radians
$\Delta\theta_2$	Spot 2 radius	0–3 radians
$kT_{\text{eff},2}$	Spot 2 effective temperature	0.011–0.5 keV
θ_{obs}	Observer inclination	1.44–1.62 radians
N_H	Neutral H column density	$0-20 \times 10^{20} \text{ cm}^{-2}$
d	Distance	$\exp[-(d - 1.136 \text{ kpc})^2 / 2(0.2 \text{ kpc})^2]$, $d \geq 1.136 \text{ kpc}$ $\exp[-(d - 1.136 \text{ kpc})^2 / 2(0.18 \text{ kpc})^2]$, $d \leq 1.136 \text{ kpc}$
A_{XMM}	XMM-Newton effective area divided by nominal area	0.9–1.1

Note. Except where noted, the prior is flat over the given range.

which the spots are disjoint or in contact, but also configurations in which they partially or completely overlap. This allows the algorithm to model approximately spots with temperature gradients, if these are favored by the data, by allowing spots with different temperatures to partially or completely overlap. This is done by assigning a number to each of the spots. For example, for the two-spot model used here, the spots are labeled 1 and 2. Points on the stellar surface that are covered by both spots are assumed to emit with the effective temperature of the lower-numbered spot. This labeling, and the precedence given to the emission by the lower-numbered spot, is not related to any other property of the spot. For example, the effective temperature of spot 1 could be either higher or lower than the effective temperature of spot 2. This approach is highly flexible, and can be generalized to any number of spots (see Miller et al. 2019 for details). For example, in addition to the best-fit solution we feature in this work, our algorithm also found a lower-likelihood solution with a large, crescent-shaped emitting region that it modeled using a spot with a temperature so low that it is essentially dark lying on top of a hotter spot. This solution is disfavored by a Bayes factor greater than 3000 as measured by MultiNest (see Section 3.5), so we have set it aside for the purposes of this work.

Section 4 of Miller et al. (2019) reported the results of an analysis of synthetic NICER pulse waveform data constructed assuming two, uniform-temperature, circular hot spots, and separately, an analysis of synthetic data assuming two, uniform-temperature, oval hot spots. In both cases, the values of the model parameters inferred from the analysis was consistent with those assumed in constructing the synthetic pulse waveform data. As another test of our analysis algorithm, in Section 4.1 we report a joint analysis of synthetic NICER and XMM-Newton pulse waveform data sets that were generated using the hot-spot model that we found provides a good description of the actual NICER and XMM-Newton data on PSR J0740+6620. This analysis includes analyses of four separate data sets, each with a realistic background.

3.3. Waveform Models

The parameters used in the pulse waveform modeling we report here are listed and defined in Table 1, with their assumed

priors. We assume that the effective temperature is uniform over a given spot.

Because PSR J0740+6620 is in a binary system, we have additional information on some of the parameters, beyond that provided by the NICER and XMM-Newton data. In particular, Fonseca et al. (2021) find that at 68% credibility the mass is $M = 2.08_{-0.069}^{+0.072} M_\odot$, the distance is $1.136_{-0.152}^{+0.174}$ kpc, and our line of sight makes an angle of $87.5 \pm 0.17^\circ$ to the orbital axis of the system.

We caution that other lines of evidence imply that PSR J0740+6620 has a lower mass. Prominent among these is that there is an observationally established correlation between the white dwarf mass and the orbital period in white dwarf binaries (Tauris & Savonije 1999; Tauris & van den Heuvel 2014), which would suggest a strong upper limit of $2 M_\odot$ to the mass of PSR J0740+6620. If the mass is indeed less than our current best estimate, the main effect will be to decrease our estimate of the radius. Although it is not generally true that the compactness $GM/(Rc^2)$ is measured to better fractional precision than the radius (see Section 4 of Bogdanov et al. 2021), this is true for PSR J0740+6620 because, as we discuss in more detail below, the most important information about the radius comes from the modulation fraction. A star that is more compact will deflect light rays coming from its surface by a larger angle, allowing a given observer to see more of the surface, which reduces the modulation fraction the observer sees. As a result, this estimate of the radius of PSR J0740+6620 is roughly proportional to its mass. Because the fractional uncertainty in the compactness of PSR J0740+6620 is large, even an estimated mass of (for example) $1.8 M_\odot$ would produce a radius posterior that largely overlaps the posterior we report here. For example, if the centroid of the mass prior were $1.8 M_\odot$ rather than $2.08 M_\odot$, and the fractional prior mass uncertainties were the same as for our standard mass priors, then our $\pm 1\sigma$ radius range would be 10.6–14.1 km rather than 12.2–16.3 km.

We use the radio observations of Fonseca et al. (2021) to impose a Gaussian prior on the mass of $M = 2.08 \pm 0.09 M_\odot$, where we have linearly added an estimated systematic error of $0.02 M_\odot$ to the mass estimate from Fonseca et al. (2021; E. Fonseca, personal communication), and an asymmetric Gaussian prior probability distribution (see Table 1 for the

functional form) centered at $d = 1.136$ kpc with 68% cumulative probability points at $d = 0.956$ kpc and $d = 1.336$ kpc, where we have placed the 68% cumulative probability points 0.03 kpc further from the central value than the quoted statistical uncertainty, to reflect the 0.03 kpc estimated systematic error (E. Fonseca, personal communication).

Our analysis depends on the inclination of our line of sight to the pulsar’s rotational axis. Millisecond pulsars such as PSR J0740+6620 are believed to be recycled by accreting matter from their companion stars (see, e.g., Bhattacharya & van den Heuvel 1991). Consequently, the orientation of their rotational axis is thought to be determined by the angular momentum of the matter accreted from the companion; accretion of this matter is expected to gradually align the pulsar’s rotational axis with the orbital axis of the system. However, this alignment is not expected to be perfect, and the current rotational axis of the pulsar may therefore be tilted by a few degrees relative to the orbital axis of the binary system. For the analyses reported here, we have adopted a prior on the angle between the observer’s line of sight and the pulsars rotation axis that is flat within 5° of the best estimate of the system’s orbital inclination, which is 87.5° , and zero outside this range. We find that the inferred mass and radius of PSR J0740+6620 are insensitive to the precise value of the orbital inclination with this range (see the corner plots in the Appendices and the discussion in Section 4.4).

3.4. Our Waveform Modeling Procedure

For PSR J0030+0451 we had only NICER data to analyze. In contrast, for PSR J0740+6620 we also had XMM-Newton imaging data from pn, MOS1, and MOS2. We therefore need to model all four data sets jointly. Once we have full models including backgrounds for all four data sets, then the log likelihood of the data given the model is simply the sum of the log likelihoods of each of the individual data sets. Thus

$$\begin{aligned} \ln \mathcal{L}_{\text{tot}} = & \ln \mathcal{L}_{\text{NICER}} + \ln \mathcal{L}_{\text{XMM-pn}} \\ & + \ln \mathcal{L}_{\text{XMM-MOS1}} + \ln \mathcal{L}_{\text{XMM-MOS2}}. \end{aligned} \quad (1)$$

In writing the total log likelihood this way, we are assuming that the data sets are uncorrelated with each other. This assumption is clearly justified for the independent NICER and XMM-Newton data. It is also reasonable for the three XMM-Newton data sets because the pn, MOS1, and MOS2 cameras are all separately mounted and read out on XMM-Newton; there is no plausible mechanism that would link counts in one detector with counts in another detector.

For each individual data set we compute the log likelihood by summing the log Poisson likelihood of the data given the model in all of the phase bins and energy channels, neglecting the log factorial term that is common to all models. That is, if we have broken the data into bins ϕ_i in rotational phase, and energy channels E_j , and in each phase-channel bin the model predicts m_{ij} counts (a positive real number) and d_{ij} counts are observed (a non-negative integer) then the Poisson likelihood of the data given the model in that phase-channel bin is $[m_{ij}^{d_{ij}}/d_{ij}!]e^{-m_{ij}}$. However, the factor $1/d_{ij}!$ is common to all models and hence can be neglected in both parameter estimation and model comparison. The log likelihood we use

is therefore

$$\ln \mathcal{L} = \sum_{\phi_i} \sum_{E_j} (d_{ij} \ln m_{ij} - m_{ij}), \quad (2)$$

where there is an implied sum over all four instruments. The NICER data have 32 phase bins. In contrast, for each of the three XMM-Newton data sets, we have only one bin in the rotational phase because the timing resolution is insufficient to divide the data further.

We now discuss the two components of the calculation of m_{ij} for each phase-channel bin in each data set: the computation of the model counts from the spots for each of the four instruments, which is straightforward, and the incorporation of phase-independent background counts, which is more involved and requires different treatments for the NICER and for the XMM-Newton data.

3.4.1. Calculation of the Model Counts from the Spots

Suppose that we are considering a particular parameter combination, i.e., we have some combination of $M, R_e, \theta_{\text{obs}}, \dots$, to assess. Given the specified observation time, this gives us a spots-only waveform that is unique up to an arbitrary definition of the time corresponding to phase 0. As described in Section 3.4 of Miller et al. (2019), once we have a candidate waveform that includes both the contribution from the spots and the contribution from the phase-independent background, we marginalize over the overall phase by assuming that near the peak the likelihood as a function of phase, $\mathcal{L}(\phi)$, can be approximated by a Gaussian in the phase.³⁴ We only need to carry out this process for the NICER data, because the XMM-Newton data have only one phase.

3.4.2. Inclusion of Unmodulated Background

By “background” we mean any X-ray counts that are not contributed by the spots. This could include space weather, optical loading, resolved or unresolved background or foreground sources, the instruments themselves, or, in the case of a binary such as PSR J0740+6620, the companion star (e.g., interactions of the neutron star particle wind with the companion star could produce X-rays). What these have in common is that those X-rays are not modulated at the rotational frequency of the star and hence with enough exposure time they contribute equally at all rotational phases.

Miller et al. (2019) therefore treated all non-stellar X-rays with a single approach, which we adopt here for the NICER data: for each energy channel independently we allow there to be a phase-independent component of any magnitude. We therefore do not assume any particular spectral form for the non-stellar emission, but in practice this results in an estimated background that varies smoothly with energy channel. The assumption of independence between channels means that, as with the phase, we fit a Gaussian to the likelihood of each energy channel as a function of added background, and can then analytically integrate the likelihood to marginalize over the background. See Miller et al. (2019) for more details. We also note (see Section 4.2 for a more thorough discussion) that future analyses may be able to incorporate estimates of the

³⁴ This is equivalent to fitting a parabola to $\ln \mathcal{L}(\phi)$ near the peak of the log likelihood. Note that the likelihood for a given phase need not be Gaussian for this approximation to be useful.

NICER background rather than letting the background be completely free.

The XMM-Newton data must be treated differently, because there is only one phase bin, in contrast to the 32 rotational phase bins that we use for the NICER data. To understand why, suppose that we treat the XMM-Newton background, for any instrument, in the same way that we treat the NICER background: each energy channel (with its single phase) can have added to it any number of background counts. Then the best model for the XMM-Newton data would have zero counts from the spots, and the background would equal the observed number of counts. This would automatically maximize the log likelihood. The NICER data are not driven to this pathological fit because there is clear modulation in the data; a phase-independent, background-only fit cannot reproduce this modulation.

Therefore, we instead assume that the blank-sky background is a Poisson realization of the full XMM-Newton background. Each of the XMM-Newton instruments has had, throughout the course of the mission, significant exposure to portions of the sky with no known X-ray sources (see Section 2). The total blank-sky count rate in the XMM-Newton energy channels of interest is $\sim 20\%$ – 30% of the count rate in our XMM-Newton observations of PSR J0740+6620, so incorporation of this background is expected to make a meaningful difference to our analysis. We note that there are other possible contributions than the blank sky to the phase-independent background. For example, as mentioned above, the interaction of the pulsar wind with the binary companion could produce X-rays. Because the total count rate is fixed, if the background is larger then the estimated total count rate from the neutron star is reduced. As a result, the modulation fraction is increased because the modulated count rate is fixed by the NICER observations. In turn, this would increase the inferred radius for PSR J0740+6620. In our treatment we effectively assume that these other contributions can be neglected.

Even with hundreds of thousands of seconds of total exposure time, the total number of counts from the blank-sky observations is too small in many of the energy channels to use Gaussian statistics. Therefore we need to use Poisson statistics. Our approach to the pn, MOS1, and MOS2 backgrounds is as follows.

Consider some particular XMM-Newton energy channel n for one of the three instruments. Suppose that the blank-sky observation had duration T_{back} and yielded b_{back} counts. If the prior on the background counts is flat, then the Poisson probability distribution for the time-averaged number of counts b_{avg} for an interval T_{back} is

$$P(b_{\text{avg}})db_{\text{avg}} = (1/b_{\text{back}}!)b_{\text{back}}^{b_{\text{avg}}} \exp(-b_{\text{avg}})db_{\text{avg}}. \quad (3)$$

Now suppose that for the XMM-Newton instrument under consideration the PSR J0740+6620 observation has a duration of T_{obs} . Then a given b_{avg} implies an expected number of counts $b_{\text{obs}} = b_{\text{avg}}(T_{\text{obs}}/T_{\text{back}})$ counts in that observation. Then because $b_{\text{avg}} = (T_{\text{back}}/T_{\text{obs}})b_{\text{obs}}$, Equation (3) implies that the normalized probability distribution for b_{obs} is

$$P(b_{\text{obs}})db_{\text{obs}} = (1/b_{\text{back}}!)[b_{\text{obs}}(T_{\text{back}}/T_{\text{obs}})]^{b_{\text{back}}} \exp \times [-b_{\text{obs}}(T_{\text{back}}/T_{\text{obs}})](T_{\text{back}}/T_{\text{obs}})db_{\text{obs}}. \quad (4)$$

If the energy channel had d observed counts in our PSR J0740+6620 observation, then the probability of

obtaining that in a model with m expected counts is

$$p(d|m) = m^d/d! \exp(-m). \quad (5)$$

If in our model we have s expected counts from the spots, then given a background b we expect a total number of counts $m = s + b$. Setting $b = b_{\text{obs}}$ as above, we find that the final likelihood of the data d given the spot counts s and the distribution $P(b_{\text{obs}})db_{\text{obs}}$ is

$$p(d|s, \text{back}) = \int_0^\infty [(s + b_{\text{obs}})^d/d!] \exp \times [-(s + b_{\text{obs}})]P(b_{\text{obs}})db_{\text{obs}}. \quad (6)$$

We find that we achieve sufficient precision for this integral by (1) setting a scale for b_{obs} that is $T_{\text{obs}}/T_{\text{back}}$ if $b_{\text{back}} = 0$, or $b_{\text{back}}(T_{\text{obs}}/T_{\text{back}})$ otherwise, and (2) performing a Simpson's rule integration from 0 to 6.8 times that scale, in intervals of 0.1 times that scale.

3.5. Parameter Estimation and Model Evaluation

As in Miller et al. (2019), we began our sampling of parameter space using the publicly available nested sampler MultiNest (Feroz et al. 2009), typically with 1000 live points, a target efficiency of 0.01 (with variable efficiency), and a tolerance of 0.1. The inverse of the efficiency parameter determines the factor by which the bounding hyperellipsoids constructed by MultiNest are expanded before samples are drawn from within them (see Section 5.2 of Feroz et al. 2009). Lower efficiency means more thorough sampling, and better accommodation of isolikelihood surfaces that are not well-described by ellipsoids. We chose an efficiency of 0.01, which is more than an order of magnitude below the standard recommendation, because we found in early tests that when we fixed the number of live points to 1000 and used efficiencies of 0.1 and 0.03 the $\pm 1\sigma$ credible regions for the radius were, respectively, 45% and 22% narrower than for our efficiency = 0.01 runs, and thus that the sampling at these higher efficiencies with 1000 live points was incomplete.

MultiNest is designed to estimate the Bayesian evidence for a model, and can also provide a quick characterization of parameter space. However, the purpose of nested samplers such as MultiNest is to compute the evidence; convergence of that calculation does not guarantee convergence of the posteriors of the parameters. Studies have shown that nested samplers of this type can give incorrectly small credible regions if isolikelihood surfaces cannot be well-represented by overlapping hyperellipsoids (Buchner 2014, 2021; Nelson et al. 2020). This appears to be the case for our PSR J0740+6620 parameter estimation; we found that the parameter space was not completely explored even at the lowest efficiencies and largest number of live points that we employed. As one example, when we used MultiNest to analyze just the NICER data with 1000 live points, a target efficiency of 0.01, and a tolerance of 0.1, the ± 1 standard deviation range in the radius posterior had 67% of the width we obtained using parallel-tempered emcee (see the next paragraph). When we used 3000 live points, a target efficiency of 0.01, and a tolerance of 0.1 the width increased, but only to 79% of the emcee width.

We therefore use a kernel density estimate of the MultiNest posteriors to produce the initial positions of the walkers, which we find expedites convergence, especially for multi-modal posteriors, in the parallel-tempered Markov chain Monte Carlo

(MCMC) sampler PT-emcee that is included in the publicly available emcee package, version 2.2.1 (Foreman-Mackey et al. 2013). Because MCMC samplers are designed to satisfy detailed balance, there is a theoretical expectation that the samples will represent the posterior. For the PT-emcee runs we had 1000 walkers per parallel-tempering temperature rung; exploratory runs with 2000 walkers per rung produced results that were consistent with what we report here. We started with five temperature rungs, but found that the number of rungs was not critical and reduced the number to two later in the sampling.

3.6. Choice of Energy Channels

In their analysis of the PSR J0030+0451 data, Miller et al. (2019) used only NICER energy channels 40 and above, because they found that inclusion of lower-energy channels biased the inferred distance. For our analysis of the PSR J0740+6620 data we use the updated 2020 July NICER response (see <https://heasarc.gsfc.nasa.gov/docs/heasarc/caldb/nicer/>). With this response there does not appear to be any bias in distance or other quantities when we use data down to energy channel 30. Analysis of the NICER data showed that the detection significance is optimized when we use channels up to 118 (see Section 2.1), but there is also an excess of counts relative to the background up to channel ~ 150 . We used a range between these (channels 30 through 123 inclusive), but exploratory runs including channels up to 150 did not produce palpably different results. Given that the width of the NICER channels is 0.001 keV, our energy channel choice corresponds approximately to photon energies between 0.3 keV and 1.23 keV. For the XMM-Newton data, we used channels 57–299 inclusive for pn, channels 20–99 inclusive for MOS1, and channels 20–99 inclusive for MOS2.

3.7. Relative Normalization of NICER and XMM-Newton

In addition to the uncertainty in the energy-dependent response of an instrument, there is inevitably some uncertainty in the overall response. That is, the overall normalization of the response is not known perfectly. A conservative upper limit on the fractional uncertainty of the relative normalization between NICER and the three XMM-Newton instruments is $\sim 10\%$.³⁵ The three XMM-Newton instruments are thought to have the same overall normalization to $\sim 2\text{--}3\%$, which are differences small enough to be neglected. Errors in the overall normalization of NICER are absorbed in the $>15\%$ half-width of our prior on the distance. However, there is still a question of whether the overall normalizations of NICER and XMM-Newton could affect our fits.

To explore this, we report in Section 4.2 an additional fit in which we add one more parameter: A_{XMM} , which is a single number that multiplies the ARFs of all three XMM-Newton instruments that we consider. We explore the consequences of allowing A_{XMM} to be between 0.9 and 1.1, i.e., the largest possible range. We therefore account for a possible inaccuracy in the ARF of NICER by the freedom in the fit to the distance to the source, and for an additional independent inaccuracy of the ARF of the XMM-Newton instruments using A_{XMM} . We

find that introduction of the parameter A_{XMM} does not produce a significant change in the posteriors of any of our parameters (for example, the radius limits are changed by at most 0.1 km) and the value of A_{XMM} that we infer is consistent with 1.

4. Results of the Analysis of the NICER and XMM-Newton Pulse Waveform Data

Our joint fit of our two-circular-spot model to the NICER and XMM-Newton data yields an equatorial circumferential radius of $R_e = 13.7_{-1.5}^{+2.6}$ km at 68% credibility for PSR J0740+6620. In this section we describe how we reach this conclusion. In Section 4.1 we demonstrate that our method produces statistically expected results when we fit our model to a synthetic data set that is drawn from a good fit to the PSR J0740+6620 data including backgrounds, and is therefore realistic. In Section 4.2 we present our fits to just the NICER data on PSR J0740+6620, and to the NICER and XMM-Newton data combined. We show that although the radius posteriors in the two analyses have significant overlap, inclusion of the XMM-Newton data favors larger radii. In Section 4.3 we show that our sampling of the parameter space has converged. In Section 4.4 we address the effects of priors, and show that other choices for the upper limit of the radius and the allowed range of the observer inclination have little effect on the -1σ radius, which is particularly important in EOS inference (see Section 5). In Section 4.5 we address the adequacy of our model. We show that our model provides acceptable fits to the phase-channel NICER data, the bolometric NICER data, and the XMM-Newton data from all three cameras. There is therefore no indication that our model for the PSR J0740+6620 data is deficient. Finally, in Section 4.6 we note some differences between our analysis method and that of Riley et al. (2021), and suggest which of those differences could be important in producing different radius posteriors between the two groups.

In this section we consider values of $R_{\text{eq}}c^2/(GM)$ up to 8.0, which corresponds to a radius ~ 24 km at $M \sim 2 M_\odot$. We find that the radius posterior does not have significant probability near this boundary, and thus that this boundary does not bias our radius estimates. When we explore the implications of results for the EOS in Section 5, we include information from nuclear theory and previous measurements.

4.1. Fits to Synthetic NICER and XMM-Newton Data

Joint analyses of NICER and XMM-Newton data have not previously been performed in the context of neutron star mass and radius estimation. It is therefore important to demonstrate that our method yields the expected results when similar synthetic data are analyzed, because for synthetic data we know the parameter values used to construct the data and can therefore judge the quality of the fit.

To generate the data we started from a good model of the PSR J0740+6620 NICER and XMM-Newton data, including all backgrounds. We then produced the synthetic data by drawing integer numbers of counts from our (in general non-integer) model expectation in each phase-channel bin for NICER, and in each energy channel for each of the three XMM-Newton cameras. We emphasize that this leads to synthetic data with fully realistic backgrounds; previous synthetic data analyses (e.g., Lo et al. 2013; Miller & Lamb 2015; Miller et al. 2019; Bogdanov et al. 2021) have instead

³⁵ For the cross-calibration as of December 2018 see the lower panel of Figure C2 in Ricci et al. (2021). For the cross-calibration as of November 2020 see slide 7 of <http://iachec.org/wp-content/presentations/2020/NICER-CrossCal-IACHEC-Markwardt-2020b.pdf> and note in particular that in the energy range relevant to our analyses, 0.3–1 keV, the flux inferred from 3C 273 agrees between NICER and XMM-Newton to better than 4%.

Table 2
Verification of Fits to Synthetic NICER and XMM-Newton Data

Sampler	$\pm 1\sigma$ (68.3%)	$\pm 2\sigma$ (95.4%)	$\pm 3\sigma$ (99.7%)
PT-emcee	6	11	12
MultiNest (Nlive = 1000, efficiency = 0.01)	4	7	9
MultiNest (Nlive = 2000, efficiency = 0.01)	4	8	11

Note. Results of fits to synthetic J0740-like NICER and XMM-Newton data using parallel-tempered emcee and two different settings of MultiNest. The right three columns show the number of parameters, out of 12, for which the values assumed in constructing the synthetic data are within the $\pm 1\sigma$, $\pm 2\sigma$, and $\pm 3\sigma$ credible regions. The PT-emcee results are consistent with statistical expectations (see Appendix A for the full posteriors). In contrast, the credible regions for both MultiNest runs are too narrow, although with more live points the sampling is more thorough.

typically used power-law backgrounds. This is therefore the most realistic synthetic data test that has been performed in the radius estimation context.

The results of our analysis of the synthetic data are summarized in Table 2, and more detail on our PT-emcee analysis is shown in Table 5 and Figure 14 in Appendix A. We see that the PT-emcee sampling yields results that are entirely consistent with statistical expectations. In contrast, MultiNest sampling with Nlive = 1000 and an efficiency of 0.01 is not adequate; for example, three of the 12 parameters have $\pm 3\sigma$ credible regions that exclude the parameter value used to construct the data. The MultiNest sampling with Nlive = 2000 and an efficiency of 0.01 is better but the credible regions are still too narrow: one of the 12 true parameter values is excluded at $>3\sigma$, and four of the 12 are excluded at $>2\sigma$. It is possible that, with more live points and/or lower efficiency, MultiNest analysis of our synthetic data would lead to the statistically expected outcomes, but results such as these form part of our motivation to use PT-emcee to determine posterior probabilities.

In summary, even with realistic backgrounds and even when analyzing jointly the NICER and XMM-Newton synthetic data, our analysis method yields statistically acceptable results.

4.2. The Mass and Radius of PSR J0740+6620

In this section we present fits to (1) just the NICER data (see Table 6 and Figure 15 in Appendix B for the best fit and uncertainties, the left-hand panels of Figure 1 for the mass and radius posteriors), (2) the NICER data and the data from the three XMM-Newton instruments, assuming the nominal calibration for XMM-Newton (see Table 7 and Figure 16 in Appendix C for the best fit and uncertainties, and the right-hand panels of Figure 1 for the mass and radius posteriors), and (3) the NICER data and the data from the three XMM-Newton instruments, with an additional energy-independent factor $0.9 < A_{\text{XMM}} < 1.1$ that multiplies the effective area of the XMM-Newton instruments (see Table 8 and Figure 17 in Appendix D for the best fit and uncertainties). The Appendices have the full corner plots for each of the three fits. In all cases we use atmosphere tables for pure hydrogen that accommodate partial ionization.

Figure 2 shows a representative spot pattern. We note that because the observer inclination is very close to the rotational equator, there is a near-symmetry between spots in the northern hemisphere (which we define to be the hemisphere of the observer) and the southern hemisphere. In addition, because the spots do not overlap, their labeling (i.e., which is spot 1 and which is spot 2) can be swapped without affecting the solution, unlike what would be the case if the spots overlapped. Thus there is an approximate four-fold degeneracy in the solution.

The inclusion of XMM-Newton data increases the radius estimates: for example, with the NICER data alone the $\pm 1\sigma$ range in radius is 10.382 km to 13.380 km, whereas when we add the XMM-Newton data the range becomes 12.209 km to 16.326 km. This is because the prime source of information about the compactness is the modulation fraction (see Miller 2016). The NICER data are dominated by background counts, which means that from that data set alone it could be that the modulation fraction is low. If other parameters such as the observer inclination and spot locations are fixed, then because light deflection increases with increasing compactness $GM/(R_e c^2)$, for a more compact star the light from spots is spread out more and thus the modulation fraction is reduced. As a result, the NICER data alone can be fit with very compact stars. But when the XMM-Newton data are included, we have a measure of the total flux from the star. This total flux is significantly less than that in the NICER data, because those data include significant background. Thus, the modulation fraction is higher than it appears to be from the NICER data alone, and hence the compactness of the star must be lower. Because the mass is known with some precision, a lower compactness implies a larger radius.

We also note that, for the same reason, if reliable information about the non-spot NICER flux were included in the NICER data then the inferred radius would increase. That is, any information that places an upper limit on the flux from the spots, whether it comes from observations of the source (in our case, using XMM-Newton) or a model of some part of the non-spot flux (which we would obtain with a background model), increases the inferred modulation fraction and thus increases the inferred radius. Hence analysis of only the NICER data on PSR J0740+6620, if augmented with NICER background data that (critically) have reliable estimates of the field-to-field variance as well as statistical uncertainties, is likely to result in an inferred radius that is close to the radius inferred when we also use XMM-Newton data. The 3C50 NICER background model of Remillard et al. (2021) is a good step in this direction, but the variance between the background in different sky pointings is not well understood.

We stress that, despite the clear shift of the radius posterior distribution to larger radii when the XMM-Newton data are included, this distribution is still broadly consistent with that obtained using the NICER data alone, even without background information. One indication of this is that there is a substantial overlap between the two $\pm 1\sigma$ ranges. To quantify this overlap we can compute, e.g., the Bhattacharyya coefficient (Bhattacharyya 1943)

$$BC(p, q) = \int \sqrt{p(x)q(x)} dx \quad (7)$$

for two normalized probability distributions $p(x)$ and $q(x)$. $BC = 1$ for identical distributions, whereas $BC = 0$ for




Figure 1. Comparison of the M – R_e corner plots obtained using two circular spots, fit to only the NICER data (left panels) and to both the NICER and the XMM-Newton data with the nominal XMM-Newton calibration (right panels). All of the results that we present use pure hydrogen model atmospheres that allow for the possibility of partial ionization. The mass prior is represented by the dashed line in each of the one-dimensional mass posterior plots. In the mass–radius plots, brighter colors indicate higher posterior probability densities; the inner contour in these plots contains 68.3% of the posterior probability whereas the outer contour contains 95.4%. In the NICER+XMM-Newton one-dimensional radius posterior plot, the dotted line shows the posterior obtained using only the NICER data. The two radius posteriors have substantial overlap, but the effect of including the XMM-Newton data is to increase the favored radius by roughly 2 km.




Figure 2. Spot locations for a representative fit, in an equal-area Mollweide projection. The horizontal bar indicates the colatitude of the observer. Because we observe from close to the rotational equator, and because the spots do not overlap, there is an approximate four-fold degeneracy in solutions.

distributions with zero overlap. In our case, $BC = 0.77$ when p is the radius posterior using only the NICER data and q is the radius posterior when the XMM-Newton data are also included, and when we represent those distributions using histograms with bin widths of 0.1 km. For context, $BC = 0.77$ for two unit-variance Gaussians whose centers are separated by roughly 1.5 times their standard deviations. This indicates that there is a significant overlap between the two distributions.

4.3. Convergence of Sampling

As we discussed in Section 3.5, we started our PT-emcee sampling of the posterior using an initial distribution of points that had previously been obtained by using MultiNest (with $N_{\text{live}} = 1000$ and efficiency = 0.01) to sample the posterior; as usual, MultiNest had run until the evidence estimate converged.

The left panel of Figure 3 shows the evolution of the -1σ , median, and $+1\sigma$ values of the radius posterior as our PT-emcee sampling of the NICER+XMM-Newton data using the nominal calibration progressed, starting with the initial parameter distribution provided by MultiNest. Each point in this plot indicates the estimates of these three radii given by a single set of 20,000 samples; the individual sets are contiguous and non-overlapping. The evolution shown was computed using approximately 1700 such sample sets; we show another view of the progressive mass–radius sampling in Figure 4. The left panel in Figure 3 shows that as the sampling by PT-emcee proceeded, the values of all three characteristic radii increased rapidly at first, indicating that the sampling achieved by MultiNest was insufficient and that the estimated values of the parameters it provided were inaccurate. Eventually the values of the three parameters ceased increasing and begin wandering up and down stochastically by small amounts, indicating that the PT-emcee estimates of the three radii had converged.

The dashed vertical line indicates the beginning of the sample sets that we used to construct our final posterior distributions (a total of 10^7 samples). The effective number of independent samples can be estimated by dividing the total number of samples by the average integrated autocorrelation time. Using typical methods for calculating the autocorrelation time for each walker (e.g., Sokal 1997) yields an estimate of ~ 4 iterations, suggesting a total of $\sim 2.5 \times 10^6$ independent samples. From sample set to sample set in the converged part of the evolution, the median, $\pm 1\sigma$, and $\pm 2\sigma$ radii vary with a standard deviation of only a few hundredths of a kilometer.

The right panel of Figure 3 shows the evolution of the three radii in our PT-emcee sampling of the model for the NICER+XMM-Newton data when a calibration parameter is included. Again we show the three radius values computed using




Figure 3. Evolution of the -1σ , median, and $+1\sigma$ values of the radius posterior as our PT-emcee sampling of the NICER+XMM-Newton data progressed. Each point in the plots indicates the estimated values of these three radii given by a single set of 20,000 contiguous, non-overlapping samples. The number of sets included in each estimate is shown by the numbers on the horizontal axis. Left: evolution of these three radii as sampling proceeds, starting from the final posterior distribution provided by MultiNest, and assuming the nominal ratio of the NICER overall effective area relative to that of XMM-Newton. As sampling proceeded, all three radii increased rapidly and then converged to values that are larger, with a broader spread, than was obtained from the MultiNest sampling. The sample sets to the right of the dashed vertical line were used to construct the final posterior distribution from this run (a total of 10^7 samples), which is what we used in our parameter estimation. Right: evolution of the three radii as sampling proceeds, starting with a kernel density estimate of the posterior distribution obtained at the end of the previous run, but now including the overall effective area A_{XMM} of XMM-Newton as a parameter in the fit. The change in the treatment of the effective area produced a small shift in the values of these three radii relative to their final values in the left panel. As sampling progressed, the three radii rapidly approached the values they had at the end of the run shown in the left panel. The 10^7 samples to the right of the dashed vertical line were used to construct the posterior distribution for this run. See the main text for more details.

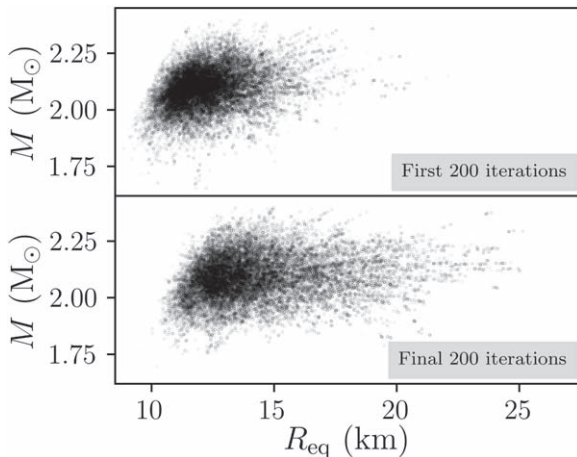


Figure 4. Another view of the evolution of the sampling in our analysis of the NICER and XMM-Newton data on PSR J0740+6620. The top panel shows the positions of the walkers in mass–radius space during the first 200 iterations of our PT-emcee sampling; these points were drawn from the posterior of an $N_{\text{live}} = 1000$, efficiency = 0.01 MultiNest run, using a kernel density estimate. The bottom panel shows the positions of the walkers during the final 200 iterations, after the PT-emcee sampling had converged. These panels correspond to the beginning and end of the left panel of Figure 3, and demonstrate the difference in the sampling between the two stages.

contiguous and non-overlapping sets of 20,000 samples. This sampling was initiated using kernel density estimation based on 10% of the final posterior samples from the PT-emcee analysis of NICER+XMM-Newton data performed using the nominal XMM-Newton calibration. As a result, the initial values of the three radii did not match perfectly the final values of the three radii in the left panel, indicating that the initial posterior distribution in the right panel does not quite match the final posterior distribution from the run shown in the left panel. However, as Figure 3 shows, these three radius estimates rapidly approached the values they had at the end of the run shown in the left panel, indicating that as the PT-emcee

sampling progressed, the posterior distribution converged quickly to the final posterior distribution in the run shown in the left panel.

We note that both our MultiNest and our PT-emcee analyses of only the NICER data, and of the NICER+XMM-Newton data that were performed using the nominal XMM-Newton calibration, assumed a mass prior of $2.11 \pm 0.09 M_{\odot}$, consistent with a preliminary radio-based mass measurement. The posterior function that was sampled to produce the three radius values plotted in the left panel of Figure 3 was computed assuming this prior. However, the posterior function that was sampled to produce the three radius values plotted in the right panel assumed a different mass prior of $2.08 \pm 0.09 M_{\odot}$, consistent with the current radio-based mass measurement. As a result, the radius values plotted in the right panel of the figure differ by about 1.5% from the radius values plotted in the left panel. The numbers in the tables and figures that summarize the posteriors produced by our analyses of the NICER-only and the NICER+XMM-Newton data assuming the nominal calibration have been re-weighted using the updated mass.

4.4. Effect of Priors

As summarized in Table 3, our posteriors are robust against several different choices of priors. In Table 3 we compare the $\pm 1\sigma$ and median radii when we analyze the NICER and XMM-Newton data using our standard priors with the $\pm 1\sigma$ and median radii when we (1) include a parameter A_{XMM} that models the relative calibration of NICER and XMM-Newton (as discussed before, our flat prior $0.9 \leq A_{\text{XMM}} \leq 1.1$ is broader than the expected uncertainty around $A_{\text{XMM}} = 1$ by a factor ~ 2), and (2) restrict our inclination prior to being flat within $\pm 0.5^\circ$ of the radio measurement, rather than our standard $\pm 5^\circ$. In each case the -1σ radius (which has especially important implications for the EOS; see Section 5) is close to our standard value. For comparison we also show in this table the $\pm 1\sigma$ radii

Table 3
Summary of Inferred Radii Using Different Priors

Prior	-1σ (km)	Median (km)	$+1\sigma$ (km)
NICER+XMM-Newton (featured result)	12.209	13.713	16.326
NICER+XMM-Newton, A_{XMM}	12.153	13.705	16.298
NICER+XMM-Newton, $\theta_{\text{obs}} = \text{radio}$	12.236	13.750	16.416
NICER only	10.382	11.512	13.380

Note. Comparison of the -1σ , median, and $+1\sigma$ radii inferred using different priors on the NICER+XMM-Newton data sets, and for the NICER-only analysis. See the text for details and note that the top, boldfaced numbers are from the analysis that we use in Section 5 for our EOS inference. When the NICER and XMM-Newton data sets are analyzed jointly the -1σ radius, which is particularly important for EOS analysis (see Section 5), is robust against different priors.

when we use only the NICER data, with no assumed knowledge of the NICER background.

4.5. Adequacy of the Models

As in Miller et al. (2019), we performed χ^2 tests to assess the adequacy of our models in describing the data. This is a one-way test: if $\chi^2 \sim \text{dof}$ for dof degrees of freedom this does not guarantee that the model is correct, but if χ^2 is sufficiently larger than dof that the probability is very low, it indicates that the model is deficient. We focus on the results of the joint fit to the NICER and XMM-Newton data using the nominal XMM-Newton calibration, given that the XMM-Newton data refine the solution substantially and that allowance for a shift in the effective area of XMM-Newton does not make a significant difference to the posterior credible regions.

In Figure 5 we show the value of χ over each of the NICER phase-channel bins we used in our analysis, for our best fit to the NICER+XMM-Newton data. There are no obvious patterns, and the χ^2/dof is 2912.37/2901 for the NICER data alone, which has a probability of 0.437 if the model is correct. In Figure 6 we compare to the data our bolometric model waveform for our best fit, and also display the residuals. Here $\chi^2/\text{dof} = 21.85/18$ (probability of 0.239) for the NICER data alone.

Our final comparison to the NICER data is in Figure 7, which compares the non-spot background in our best model (black line) with the blank-sky NICER background inferred using the 3C50 model of Remillard et al. (2021). Because the 3C50 model is for a blank sky, i.e., directions where there are no known sources, it should provide a lower limit to the background in the direction of a particular source such as PSR J0740+6620. Indeed, we see that our estimated background counts, while consistent with the 3C50 model at energy channels ≈ 80 and higher, are above the model at lower energy channels. This is likely to be due to particular sources within the NICER field of view that can be identified in the XMM-Newton image of the field (see Wolff et al. 2021).

Our model is therefore a good fit to the NICER data.

In Figure 8 we compare our best NICER+XMM-Newton model with the XMM-Newton data. Because there are few counts per energy channel, we evaluated the quality of our fit by generating 10^5 synthetic data sets by performing Poisson draws using our best model (including the blank-sky counts) in every energy channel in all three instruments. We then compared our model to the J0740 data, and placed that in the context of the comparison of our model to the synthetic data sets, in two ways.




Figure 5. Value of χ in each of the 3008 NICER phase-energy bins (32 phase bins and 94 energy channels), for the best fit of our energy-resolved waveform model to the NICER+XMM-Newton data. As is expected for a good fit, no patterns in the values of χ are evident as a function of phase or energy. The χ^2/dof is 2912.37/2901, which has a probability of 0.437 if the model is correct.

1. We computed the total log likelihood of all three XMM-Newton data sets for PSR J0740+6620 given our model. We also computed the distribution of the total log likelihoods of each synthetic data set given our model. The total log likelihood of the real data is at the 75th percentile of the log likelihoods from the synthetic set (98th percentile for the pn data alone, 17th percentile for the MOS1 data alone, and 12th percentile for the MOS2 data alone).
2. We also calculated the Kolmogorov–Smirnov (K-S) probability that the cumulative counts in our model, as a function of the energy channel, are drawn from the same distribution as the PSR J0740+6620 data. We then calculated the distribution of the K-S probabilities for the synthetic data set. We found that the K-S probability for the XMM-Newton pn data on PSR J0740+6620 is at the 10th percentile of the corresponding probabilities for the synthetic data sets, the K-S probability for the MOS1 data on PSR J0740+6620 is at the 50th percentile, and the K-S probability for the MOS2 data on PSR J0740+6620 is at the 31st percentile.

As a final check, we note that the XMM-Newton DDT data have 94 counts in the pn camera in the energy range we investigate, 59 counts in the MOS1 camera, and 67 counts in the MOS2 camera. In our best fit, we have respectively 114.88 counts (1.95 standard deviations high), 49.85 counts (1.30 standard deviations low), and 59.78 counts (0.93 standard deviations low). Neither the total count numbers nor the comparisons with synthetic data indicate a problem. Thus, at least based on these tests, our model is fully consistent with all three XMM-Newton data sets.




Figure 6. Top: comparison of the 32-phase bolometric waveform constructed using the NICER data on PSR J0740+6620 with the bolometric waveform model that best fits the NICER+XMM-Newton data. The dashed blue line shows the unmodulated background that was added to the counts produced by the two hot spots as part of the fitting procedure (see Section 3.4.2). Bottom: resulting value of χ as a function of phase. The χ^2/dof is 21.85/18, which has a probability of 0.239 if the model is correct.

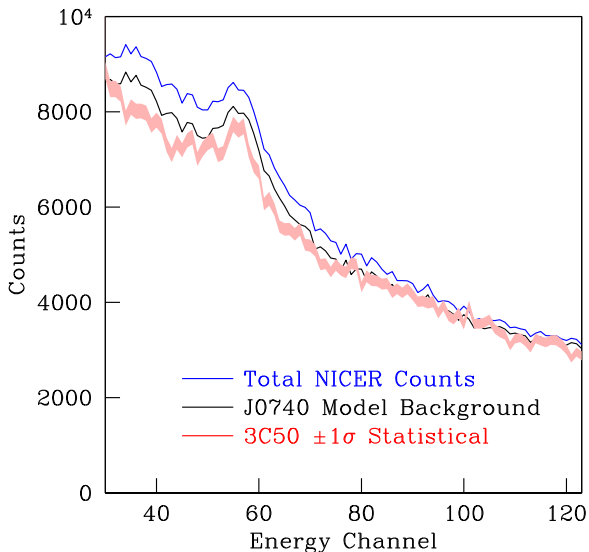


Figure 7. Inferred non-spot counts as a function of NICER energy channel in our best model (black line) compared with the $\pm 1\sigma$ range of the 3C50 background model (pink band) of Remillard et al. (2021). For comparison, we also show the total NICER counts per channel (blue line). Above channel 80 our inferred background is consistent with the 3C50 model. The excess in our inferred background at lower energies is likely due to the presence of other sources in the NICER field, which can be seen in XMM-Newton images.

4.6. Potentially Important Differences between Our Analysis and That of Riley et al. (2021)

The parallel analysis of the NICER and XMM-Newton data by Riley et al. (2021) finds a smaller credible interval for the radius estimate and, most important for the inferred constraints on the EOS, a value for the radius at the -1σ contour that is ≈ 0.8 km smaller than the value we report here. There are a number of differences between their analysis procedure and ours. For example, Riley et al. (2021) assume that the inclination of the observer’s line of sight to the pulsar’s

rotational axis is identical to the measured inclination of the observer’s line of sight to the orbital axis of the binary system, whereas we allow the inclination of the observer’s line of sight to the pulsar’s rotational axis to differ by as much as $\pm 5^\circ$ from the measured inclination of the observer’s line of sight to the orbital axis. However, as we have discussed in Section 4.4, this difference between the priors on the inclination of the observer’s line of sight to the pulsar’s rotational axis used in the two analyses does not produce a significant difference in the value of the radius at the -1σ credibility contour. Another unimportant difference is the ionization fraction of our hydrogen atmospheres: we use tables that allow for the possibility of partial ionization, whereas Riley et al. (2021) used atmospheres with fully ionized hydrogen. As we indicate in Section 3.1, we found negligible changes in the radius posteriors when we used fully ionized hydrogen atmospheres.

There appear to be four differences in the analysis procedures used by the two groups that could contribute to the difference of ≈ 0.8 km in the values of the radius at the -1σ credibility contour that the two groups report in their headline results ($R = 12.2$ km at -1σ reported here and $R = 11.4$ km at -1σ reported by Riley et al. 2021).

1. We allow the relative calibration between NICER and XMM-Newton to deviate by up to $\pm 10\%$ from its nominal value, a deviation that is a factor ~ 2 times larger than the measured deviation (again see slide 7 of <http://iachec.org/wp-content/presentations/2020/NICER-CrossCal-IACHEC-Markwardt-2020b.pdf> and focus in particular on the energy range 0.3–1 keV that is relevant for the present analysis; here the deviation is $< 4\%$). In contrast, Riley et al. (2021) allow the relative calibration to deviate by several tens of percent, which is many times larger than the measured deviation; for example, in their best fit, A_{XMM} is effectively equal to 0.574. Elsewhere in their paper, Riley et al. show that when they restrict the relative calibration to a range consistent




Figure 8. Comparison of the data from the three XMM-Newton cameras (labeled at the top of each plot) with our best NICER+XMM-Newton model. In each plot the black histogram is the data, the dotted blue line is the blank-sky background scaled to the duration of the observations, and the red line is the best model including the background. In each case the energy channel range is the range we use in our analysis. As described in the text, our model provides a good fit to the data.

with the measured range, they find $R = 11.75$ km at the -1σ contour. Thus, when they require the relative calibration between NICER and XMM-Newton to have a value consistent with the measured values, the difference between their value of the radius estimate at the -1σ contour and our value is reduced by almost a factor of two.

- When estimating the radius of PSR J0740+6620, we allowed $R_{\text{eq}}c^2/(GM)$ to be as large as 8.0 (corresponding to $R_{\text{eq}} \approx 24.5$ km at $M = 2.08 M_{\odot}$) and found that the posterior radius probability density is negligible near this boundary. Hence, this choice for the boundary of our search volume did not affect our estimate of the radius of PSR J0740+6620. In contrast, when Riley et al. estimated the radius of PSR J0740+6620, they sampled the radius parameter space only at radius values < 16 km. To investigate the results we would have found if we had sampled the radius parameter space only at $R < 16$ km as they did, we required $R < 16$ km in a test analysis and found $R = 12.06$ km, 13.28 km, and 14.82 km at the -1σ contour, the median value, and the $+1\sigma$ contour. Even when we imposed $R < 16$ km and also allowed the relative calibration factor A_{XMM} to vary freely in the range $0.9 < A_{\text{XMM}} < 1.1$, a factor > 2 deviation larger than its measured value, we found that this had only a very small effect on the position and width of the credible interval of the radius, which changed only slightly to $R = 12.00$ km, 13.28 km, and 14.80 km at the -1σ contour, the median value, and the $+1\sigma$ contour. Thus, these two differences between our analysis procedure and the analysis procedure used by Riley et al. (the different

breadths of the allowed deviation of the relative calibration of the NICER and XMM-Newton instruments and artificially limiting or not limiting the sampling of the radius parameter volume) account for ≈ 0.55 km of the 0.8 km difference between the value of the radius at the -1σ contour reported by Riley et al. and the value of the radius at the -1σ contour that we found.

- We used PT-emcee to thoroughly sample the multi-dimensional posterior, using MultiNest only to create an initial rough sample that we then used to speed up the sampling using PT-emcee. In our analysis of synthetic pulse waveform data constructed to mimic the actual NICER and XMM-Newton data (see Section 4.1), we found that when we used only MultiNest, even with values of the MultiNest parameters chosen to try to achieve thorough sampling, the credible regions it generated were too small, i.e., they too often failed to include the values of the pulse waveform parameters that had been used to generate the synthetic pulse waveform data. In contrast, when we used PT-emcee we found credible regions that were consistent with the values of the parameters chosen when generating the synthetic data (see Table 2 for more details). Figure 3 shows the detailed evolution of the -1σ , median, and $+1\sigma$ radii throughout the course of our PT-emcee sampling. Riley et al. used only MultiNest for all of their analyses and found that their posteriors continued to broaden with more thorough sampling for all of the live point numbers that they explored.

4. We treated the distribution of the blank-sky counts in each energy channel as a realization of a Poisson distribution for the observed number of counts. Riley et al. instead assumed a flat prior for the distribution of the blank-sky counts in each energy channel, centered on the flux implied by the observed number of counts with a variety of widths equal to various multiples of the standard deviation of the observed number of counts.

5. Implications for the EOS

Our aim is to combine multiple lines of data relevant to the properties of cold, catalyzed, dense matter, using a Bayesian procedure that is both statistically rigorous and computationally practical. In this section we divide our approach into two subtasks: (1) likelihood computation and parameter estimation for a given EOS family (Section 5.1), and (2) different approaches to generating candidate EOSs (Section 5.2). Finally, we show in Section 5.3 that with the inclusion of the mass and radius estimates for PSR J0030+0451 and PSR J0740+6620, the constraints on the EOS depend relatively weakly on the EOS parameterization assumed in the analysis, up to a few times nuclear saturation density.

5.1. Statistical Method

There is a growing number of papers on EOS inference from astronomical and/or laboratory data (e.g., Agathos et al. 2015; Lackey & Wade 2015; Alvarez-Castillo et al. 2016; Margalit & Metzger 2017; Annala et al. 2018; Most et al. 2018; Radice et al. 2018; Rezzolla et al. 2018; Riley et al. 2018; Tews et al. 2018; Greif et al. 2019; Kiuchi et al. 2019; Landry & Essick 2019; Shibata et al. 2019; Capano et al. 2020; Cierniak & Blaschke 2020; Dietrich et al. 2020; Landry et al. 2020; Li & Magno 2020; Raaijmakers et al. 2020; Pang et al. 2020; Zhang & Li 2020; Essick et al. 2020b; Xie & Li 2021). Here we follow the procedure described in Miller et al. (2020). Their procedure is fast and flexible enough to incorporate multiple types of data.

To elaborate, suppose that we have i different types of measurements, which might include laboratory experiments, or the measurement of a high mass from a pulsar, or the measurement of a tidal deformability, a radius, a moment of inertia, and so on. Suppose also that for measurement type i we have $j = 1, 2, \dots, j(i)$ independent measurements of that type; these could be different measurements of the same source, or measurements of different sources.³⁶ We wish to evaluate a set k of EOSs; in the next section we discuss the selection of k . If the likelihood of data set (i,j) given EOS k is $\mathcal{L}_k(i, j)$, then the likelihood of all of the data sets given EOS k is simply the product of the likelihoods:

$$\mathcal{L}_k = \prod_i \left[\prod_{j=1}^{j(i)} \mathcal{L}_k(i, j) \right]. \quad (8)$$

Note that in some cases one may need to marginalize over extra parameters. As discussed in Section 2 of Miller et al. (2020),

³⁶ For example, here we present our radius measurements, which are based on X-ray data, using a prior on the mass given by radio observations. This prior, however, is obtained by starting with a flat mass prior and then updating the mass distribution using the radio data. Thus our procedure is equivalent to starting with a flat mass prior and then computing the joint likelihood of the radio and X-ray data given candidate masses and radii.

there are some parameters that are expected to be the same for all neutron stars (such as the parameters that describe the EOS beyond nuclear saturation density, given that we expect the EOS to be unique at such densities), some that can differ between neutron stars but are the same for any observation of a given neutron star (such as the distance, observer inclination, stellar mass and radius, and central density), and some that could vary between different observations of a given neutron star (such as the absorbing column if the star is in a binary and there is a significant wind). It would be incorrect to marginalize over parameters separately for different data sets that share those parameters; we do not perform such marginalizations in Equation (8). In contrast, marginalization can be performed for a likelihood applying to a given data set, over parameters that apply only to that data set. For example, to compute the likelihood of a NICER posterior in mass and radius for a given neutron star it is necessary to integrate over the prior for the central density, which applies to that star and is not the same for all EOSs because different EOSs have different maximum stable central densities.

Given the calculation of the likelihood \mathcal{L}_k of the full set of data assuming a specific EOS k (see Miller et al. 2020 for more details), then as usual in a Bayesian analysis the posterior probability P_k of the EOS is proportional to the prior probability q_k of the EOS times the likelihood:

$$P_k \propto q_k \mathcal{L}_k. \quad (9)$$

We could use this procedure to compare individual EOSs from the literature. However, our primary goal is to produce a posterior probability distribution for, say, the pressure P as a function of the total energy density ϵ or of the number density n of baryons. To do this, we simply calculate the pressure at specified ϵ or n for each EOS k , weight it by \mathcal{L}_k , and then sum over all k to find the final distribution of P at a given ϵ or n .

For each specific EOS selected from a given EOS family we need to compute the likelihood of all relevant data sets for that EOS. We largely follow the procedure of Miller et al. (2019) in our treatment of our data sets. In particular, we apply kernel density estimation (other methods exist, e.g., random forest regressors; see Hernandez Vivanco et al. 2019) to produce an estimated continuous probability distribution from our discrete samples. When we do so, we use the standard bandwidth from Silverman (1986) for our samples in mass and tidal deformability, but multiply this bandwidth by 0.1 for our mass–radius samples because we found previously that the standard bandwidth significantly broadened the radius posterior (see Miller et al. 2019 for more details, and Essick et al. 2020a for a more detailed approach to selecting the bandwidth). We modify the procedure of Miller et al. (2020) in two ways.

1. In Miller et al. (2019, 2020) the prior on the central density was flat between the density that produces a $1 M_\odot$ star and the density that produces the maximum-mass non-rotating neutron star for the EOS under consideration. Because the central density changes rapidly near the maximum mass, this gives relatively greater prior weight to higher-mass stars. Here our prior on the central density is instead quadratic in the central density between the $1 M_\odot$ density, ρ_{\min} and the maximum-mass density ρ_{\max} . That is, we select uniformly in $0 \leq x \leq 1$, where the central density is $\rho_c = \rho_{\min} + x^2(\rho_{\max} - \rho_{\min})$. This

produces a more even prior distribution on the masses. However, we assign zero prior weight to any densities between ρ_{\min} and ρ_{\max} that produce unstable stars; in this way we explicitly treat EOSs that have two stable ranges of central density separated by an unstable range of central density (these are sometimes called “twin stars”), when they arise. In practice, the different prior distribution for the central density does not change the EOS posteriors significantly compared with the approach in Miller et al. (2019).

2. Because we include the high-mass binary merger GW190425 (Abbott et al. 2020a) in our analysis, and the total mass of the binary is high enough that one of the compact objects might have been a black hole, we treat tidal deformabilities differently than in Miller et al. (2019, 2020) (although we also note that the large distance to and high mass of GW190425 mean that it adds only a small amount of information). Our approach is to select the central density of the lower-mass star, up to the central density such that the implied mass of the other star equals that of the lower-mass star. Given the mass of the lower-mass star, and given that the chirp mass (equal to $(m_1 m_2)^{3/5} / (m_1 + m_2)^{1/5}$ for stellar masses m_1 and m_2 in a binary) is known with high precision, we know the mass of the higher-mass object to the precision that we know the chirp mass. We marginalize over the chirp mass as in Equation (15) of Miller et al. (2020). If for the EOS and central density under consideration the higher-mass object is a neutron star, then we compute the tidal deformabilities of both stars, using the same EOS, following the prescription of Hinderer (2008; see the erratum at Hinderer 2009). If instead for the EOS and central density under consideration the higher-mass object is a black hole, then we assume that its tidal deformability is zero (Binnington & Poisson 2009; Damour & Nagar 2009). Note that for rotating black holes the tidal deformability may not be exactly zero (see Chia 2020; Le Tiec & Casals 2021; Le Tiec et al. 2021 for an ongoing discussion), but it is small enough that the difference from zero is not important for our calculations.

To implement our analysis, we need a mechanism to select a set of EOSs. We now discuss our three different approaches to this selection.

5.2. EOS Models

There exist many different nuclear physics computations of the EOS of cold, catalyzed matter both below and above nuclear saturation density. It is usually agreed that our understanding of matter below saturation density is fairly good, based on the idea that we can perform laboratory experiments that probe this density regime. One caveat is in order: the matter that can be examined in laboratories has roughly equal numbers of neutrons and protons (the most asymmetric stable nuclei, such as ^{208}Pb , have roughly 50% more neutrons than protons). In contrast, neutron star matter in the vicinity of saturation density is more than 90% neutrons. Thus although from the standpoint of density the regime is familiar, theoretical arguments not directly supported by experiment are necessary to extrapolate to very asymmetric matter.

Nonetheless, we will follow standard procedure and assume that we know the EOS up to a threshold density, which we take to be half of saturation density because this is approximately the crust–core transition density (Hebeler et al. 2013). Note that some calculations suggest that different treatments of the crustal EOS could introduce radius uncertainties as large as 0.3 km (Fortin et al. 2016; Gamba et al. 2020). Below half of nuclear saturation density we use the QHC19 EOS (Baym et al. 2019). Above half of nuclear saturation density, we then need to model the EOS, where there are well motivated and carefully constructed models, but where these models have not been validated by laboratory measurements.

A necessary output of our EOS analysis is an understanding of how strongly our conclusions depend on the type of EOS model that we employ. Therefore, in the rest of this section we discuss the three EOS model families that we use in our analysis. First we present two parameterizations which have been widely used in the literature. We then turn our attention to a newer method using Gaussian processes (GPs) that was introduced in this context by Landry & Essick (2019).

5.2.1. Parametric Models

The most common approach to EOS inference has used parameterized models. That is, we assume that there is some set of parameters α that completely describe the EOS at high densities. Given priors on α , we can generate a large number of sample EOSs for which we then compute the likelihood of the assembled data. As discussed above, this allows us to compute posteriors for $P(\epsilon)$ or $P(n)$. The two parameterizations we use are as follows.

Piecewise polytrope. In this model, above a number density $n = n_s/2$ the pressure is given by a series of polytropes: $P = K_1 n^{\Gamma_1}$ for $n_s/2 \leq n \leq n_2$, $P = K_2 n^{\Gamma_2}$ for $n_2 \leq n \leq n_3$, and so on, where the coefficients K_i are chosen to ensure continuity of pressure between density segments. Variants of this model differ in the number of segments, the priors on the polytropic indices Γ_i , and whether the transition densities n_1, n_2, \dots are fixed or can vary. We use the following assumptions (compare with Miller et al. 2019, which used the same assumptions except for restricting the first polytropic index to the range [2, 3]): the polytropic index from $n_s/2$ to $n_2 \in [3/4, 5/4]n_s$ is $\Gamma_1 \in [0, 5]$; from n_2 to $n_3 \in [3/2, 5/2]n_s$ is $\Gamma_2 \in [0, 5]$; from n_3 to $n_4 \in [3, 5]n_s$ is $\Gamma_3 \in [0, 5]$; from n_4 to $n_5 \in [6, 10]n_s$ is $\Gamma_4 \in [0, 5]$; and from n_5 to ∞ is $\Gamma_4 \in [0, 5]$ (all priors are uniform in the listed range). When Γ is close to zero, the pressure is nearly independent of the density, which is what one expects near a phase transition and which can produce the twin stars discussed earlier. If at some density below n_5 the implied adiabatic sound speed $c_s = (dP/d\epsilon)^{1/2} > c$, the speed of light, then we set $dP/d\epsilon = c^2$ at that density.

Spectral parameterization. A different parameterization was introduced by Lindblom (2010, 2018). In this approach, the polytropic index is a continuous function of the pressure, $\Gamma(P)$:

$$\Gamma(P) = \exp\left(\sum_i \gamma_i x^i\right). \quad (10)$$

Here $x \equiv \log(P/P_0)$ and P_0 is the pressure at $n_s/2$. We follow Abbott et al. (2018), Carney et al. (2018), and Miller et al. (2019) in expanding to $O(x^3)$, and in using uniform priors $\gamma_0 \in [0.2, 2]$, $\gamma_1 \in [-1.6, 1.7]$, $\gamma_2 \in [-0.6, 0.6]$, and $\gamma_3 \in [-0.02, 0.02]$. As with the piecewise polytrope, if $c_s > c$ at

some density then we set $c_s = c$ at that density. Unlike the piecewise-polytropic parameterization, the spectral parameterization with these ranges of γ_i cannot easily simulate a phase transition. However, in other ways it is more flexible, with fewer parameters, than the piecewise-polytropic model.

5.2.2. Models Using Gaussian Processes

A newer approach to EOS modeling using GPs was introduced in Landry & Essick (2019) (see also the subsequent work in Landry et al. 2020; Essick et al. 2020a, 2020b). Because this method is currently less commonly used in the community than the parametric approach, we will summarize it in more detail.

The essence of GP estimation of a function f (see Rasmussen & Williams 2006 for an excellent introduction) is that we would like to have the outcome of our analysis be the joint probability density for the function values $f(\mathbf{x})$ at inputs \mathbf{x} (which we can represent as x_i), which we assume to be a multivariate Gaussian distribution with means $\boldsymbol{\mu} = \mu_i$ and covariance matrix $\Sigma = \Sigma_{ij}$. That is, we would ultimately like to obtain

$$f(\mathbf{x}) \sim \mathcal{N}(\boldsymbol{\mu}, \Sigma), \quad (11)$$

where \mathcal{N} indicates a normal distribution.

A key assumption made in GP estimation is that correlations between the function values f_i, f_j at inputs x_i, x_j can be represented using a covariance kernel K that is a function of x_i and x_j . That is, if our joint distribution is

$$f_i | x_i \sim \mathcal{N}(\langle f_i \rangle, \Sigma(f_i, f_j)), \quad (12)$$

where $\langle f_i \rangle$ is the expectation value for f at x_i and $\Sigma(f_i, f_j)$ is the covariance matrix, then we assume we can write $\Sigma(f_i, f_j) = K(x_i, x_j)$.

It is reasonable to assume that when x_i and x_j are closer to each other, K is larger. Given that covariance matrices are symmetric, it also must be that $K(x_i, x_j) = K(x_j, x_i)$ for any x_i and x_j . A common and much more drastic simplification is to assume that K depends on only the distance between the points: $K(x_i, x_j) = K(|x_i - x_j|)$. We use a Gaussian kernel, which in this context is called a ‘‘squared-exponential kernel’’:

$$K_{\text{se}}(x, x') = \sigma^2 \exp\left(-\frac{(x - x')^2}{2\ell^2}\right), \quad (13)$$

which has the properties listed above. Here the hyperparameter σ determines the strength of the overall correlation and the hyperparameter ℓ gives the correlation length scale. For example, $\ell \rightarrow 0$ would mean that all of the points are independent of each other, and $\sigma \rightarrow 0$ would mean that there is negligible variance.

Transformation of dependent variable. There is one additional important qualitative aspect of the analysis of EOS using GPs. A true Gaussian has a domain of $-\infty$ to $+\infty$. However, if we think about an EOS as $\epsilon(P)$, then because ϵ is non-negative it cannot range from $-\infty$ to $+\infty$. We could instead use $\log(\epsilon)$ as a function of $\log(P)$, because if we express ϵ in some units, e.g., cgs units, $\log(\epsilon)$ can range from $-\infty$ to $+\infty$. However, this approach is also sub-optimal because it could lead to unstable ($dp/d\epsilon < 0$) or acausal ($dp/d\epsilon > c^2$) EOS. In principle one could simply discard such EOSs when they are drawn from the GP, but this would be inefficient.

Therefore, we follow Lindblom (2010) and Landry & Essick (2019) in defining a new variable:

$$\phi \equiv \ln\left(c^2 \frac{d\epsilon}{dP} - 1\right). \quad (14)$$

Because the adiabatic speed of sound c_s is given by $c_s^2 = dP/d\epsilon$, at the boundary of thermodynamic stability ($c_s^2 \rightarrow 0$), $\phi \rightarrow +\infty$, and at the boundary of causality ($c_s^2 \rightarrow c^2$), $\phi \rightarrow -\infty$. Thus if we construct a GP in $\phi(\log P)$, all values of ϕ correspond to stable and causal EOSs.

The next choice we need to make is how closely we wish to follow tabulated EOSs in our function $\phi(\log P)$. Landry & Essick (2019) define two approximate limits: the ‘‘model-informed’’ EOS prior, in which the GP is strongly conditioned on a set of EOSs proposed in the literature, and the ‘‘model-agnostic’’ EOS prior, in which the GP is only loosely related to existing EOS proposals. The model-agnostic approach is much less computationally demanding and is less biased by existing EOSs, so it is the path that we follow.

We then need to choose the parameters that describe our GP. Landry & Essick (2019) note that it is advantageous to look for approximate trends in the function we wish to model (here $\phi(\log P)$) and then to produce a GP for the residuals. We find from sets of tabulated EOS (e.g., those at the CompOSE website <https://compose.obspm.fr/table/families/3/>) that in the log pressure range of interest ($\log_{10} P$ (erg cm⁻³) ~ 32 – 36) the trend is roughly linear. We use $\phi_0 = 5.5 - 2.0(\log_{10} P - 32.7)$ as our approximate trend, but construction of the GP does not require an exact fit. For our kernel, we use the squared-exponential kernel of Equation (13) and, based on the spread between the EOS above in the $\log_{10} P$ (erg cm⁻³) ~ 32 – 36 range, we choose $\sigma = 1$ and $\ell = 1$.

We note that the GP framework allows for substantial flexibility beyond our particular choice (see, e.g., Essick et al. 2020a). For example, although $\sigma = 1$ and $\ell = 1$ matches well most existing EOSs in the density range that has a discernible impact on the maximum mass, radius, and tidal deformability of neutron stars, when combined with our $\phi_0 = 5.5 - 2.0(\log_{10} P - 32.7)$ trend it leads to sound speeds approaching the speed of light at densities several times nuclear saturation density. A different choice of hyperparameters, e.g., a larger σ , would not necessarily require such high sound speeds; see for example Figure 2 of Landry et al. (2020).

5.3. EOS Results

The constraints we obtain on the EOS at different densities depend both on the EOS model and priors and on the data sets that we include. Here we present the EOS constraints for each of our three models, with progressive incorporation of more data. We also plot the distributions of mass versus radius and the maximum mass for non-rotating neutron stars for different amounts of data and using the GP model. For a given EOS and central density, we solve the Tolman–Oppenheimer–Volkoff equation (Oppenheimer & Volkoff 1939; Tolman 1939) to obtain the mass and radius. Thus for the purpose of calculating stellar structure we implicitly treat the star as non-rotating. This is an excellent approximation at the 346.53 Hz rotation frequency of this pulsar (Cromartie et al. 2020), because this is only $\sim 20\%$ of the mass-shedding frequency and thus the expected deviation from the structure of a non-rotating star is much smaller than our measurement precision (e.g., see




Figure 9. Comparison of the pressure ranges as a function of number density for the three EOS models described in the text, as a function of the data sets included in the analysis. For each line type and color in each panel, the lower line shows the 5th percentile, and the upper line shows the 95th percentile, of the pressure as a function of density. Colors indicate the EOS model used in the analysis: black for a piecewise-polytrope model (see the first part of Section 5.2.1), blue for the spectral model (see the second part of Section 5.2.1), and red for a model based on Gaussian processes (see Section 5.2.2). We show the results for our priors (dotted lines in each panel); our priors plus symmetry energy measurements, the existence of three high-mass pulsars, and two tidal deformability upper limits (dashed-dotted lines in each panel); those plus the mass–radius posteriors on PSR J0030+0451 from Miller et al. (2019) (dashed lines in each panel); and finally all of those plus the mass–radius measurement of PSR J0740+6620 that we report here (solid lines in each panel). The lower right panel shows the full set of constraints, including our PSR J0740+6620 radius measurement, for each of our three EOS families. Inclusion of the masses and radii of PSR J0030+0451 and PSR J0740+6620 tightens significantly the EOS models in the vicinity of $n/n_s \sim 1.5$ – 5 [$\log_{10}(n/n_s) \sim 0.2$ – 0.7]. This indicates that with the NICER and XMM-Newton data added, in this density range the posterior is now dominated by the data rather than by the priors.

Figure 13 of Miller et al. 2019; at this mass and rotation frequency, the radius increase is at most 0.2 km compared with a non-rotating star).

In Figure 9 we show, for each EOS model, the middle 90% range for the pressure as a function of number density when we incorporate successively more constraining measurements. The upper left panel shows the results using the piecewise-polytropic model, the upper right panel shows the results using the spectral model, and the lower left panel shows the results using the GP model. Our first constraints (dotted lines) use only the prior for each EOS model. The next constraints include a Gaussian prior of $S = 32 \pm 2$ MeV for the symmetry energy at nuclear saturation density (see Tsang et al. 2012),³⁷ the existence of three high-mass pulsars (Antoniadis et al. 2013; Arzoumanian et al. 2018; Cromartie et al. 2020), and the tidal deformability posteriors from gravitational wave observations

of GW170817 (Abbott et al. 2017, 2018; De et al. 2018) and GW190425 (Abbott et al. 2020a). Other constraints on the symmetry energy are possible; for example, Drischler et al. (2020) find a tighter range for the symmetry energy of $S = 31.7 \pm 1.1$ MeV and analyses of PREX-II data (Adhikari et al. 2021; Reed et al. 2021; Yue et al. 2021) suggest a higher symmetry energy of $S \approx 38 \pm 4$ MeV. The third set of constraints also includes the mass–radius posteriors from the Miller et al. (2019) analysis of NICER data on PSR J0030+0451. Finally, the last set of constraints also includes our mass–radius posteriors for PSR J0740+6620. The lower right panel collects the final set of constraints for all three EOS models. In each case, we only display the results for our fit to both the NICER and the XMM-Newton data using the nominal XMM-Newton calibration. When we include our mass–radius posteriors for PSR J0740+6620, we use only the masses of PSR 1614–2230 and PSR J0348+0432 in our high-pulsar-mass data, i.e., we do not double-count the high mass of PSR J0740+6620.

Several trends are evident in this figure.

1. The 5%–95% pressure range at densities $n \approx (1.5 - 5)n_s$, after incorporation of all measurements, has been tightened substantially compared with the prior range and now depends relatively weakly on the EOS model. For example,

³⁷ The symmetry energy S is the difference between the total energy per nucleon of pure neutron matter and E_{sym} , the energy per nucleon of symmetric nuclear matter. Here we approximate S using $\epsilon/n - m_n c^2$. Given that the matter at nuclear saturation density is not pure neutrons, $\epsilon/n - m_n c^2$ could be a few percent different from S (see Latimer & Prakash 2016 for a discussion). However, the inclusion of an S constraint has only a minor effect; for example, with our constraint, the full $\pm 1\sigma$ range of radii at $M = 1.4 M_\odot$ over all three EOS frameworks is 11.8–13.1 km, whereas without any S constraint the range is 11.6–13.0 km.

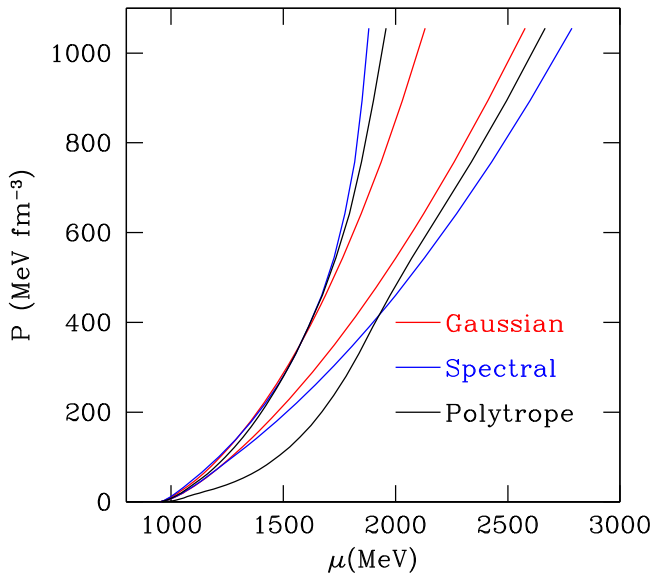


Figure 10. The 5th to 95th percentile range of the pressure as a function of the baryon chemical potential, including the neutron rest mass-energy, for all three of our models of the EOS: the red line is for the Gaussian process EOS, the blue line is for the spectral EOS, and the black line is for the piecewise-polytropic EOS. Here we include all measurements: symmetry energy, high-mass pulsars, tidal deformability from gravitational wave measurements, and measurements of the mass and radius of PSR J0030+0451 and PSR J0740+6620.

using the priors the maximum difference in the 5th percentile of the log pressure in this density range, among all three EOS frameworks, is 0.61; after incorporation of all data, the maximum difference has dropped to 0.34. For the 95th percentile, the maximum difference in the log pressure among all our EOS frameworks in this density range has dropped from 0.86 to 0.21. This is encouraging, because it means that for this density range the posterior is now dominated by the data rather than by the priors.

2. This convergence of constraints requires the NICER data on PSR J0030+0451 and the NICER+XMM-Newton data on PSR J0740+6620. Without these data, the pressure ranges differ significantly among the three models.
3. As expected, there is much less convergence at higher densities ($\gtrsim 5n_s$) and at lower densities ($\lesssim 1.5n_s$). This is because low and high densities contribute less to the radius, mass, and tidal deformability of neutron stars. See Lattimer & Prakash (2007, 2016), Xie & Li (2020), and Drischler et al. (2021) for the detailed theoretical context.
4. Finally, we note that the constraining power of our measurement of PSR J0740+6620 comes entirely from the strong lower limit on its radius. Other measurements, particularly the non-detection of tidal deformability from GW170817, already strongly exclude radii as large as even our $+1\sigma$ value of >16 km. However, the evidence that the -1σ radius is $\gtrsim 12$ km, and that the -2σ radius is $\gtrsim 11$ km, adds significantly to our understanding of dense matter.

In Figure 10 we present another view of the EOS, after all measurements are incorporated. Here we show the pressure versus the baryon chemical potential $\mu = (\epsilon + P)/n$ (including the baryon rest mass-energy) for all three of our EOS models.

In Figure 11 we present the 5th to 95th percentile range, and the 25th to 75th percentile range, of the radius as a function of

mass, using our GP model, for the same set of measurements as in Figure 9. In this figure we also plot (faint lines) representative mass–radius relations drawn from our set of EOSs, where in each panel the probability of drawing an individual EOS is proportional to its statistical weight given the measurements considered in that panel. Here we only plot masses below the 95th percentile of the maximum gravitational mass of a non-rotating neutron star, given the measurements being considered. The radius and mass of PSR J0030+0451, and our current measurement of the radius of PSR J0740+6620, tighten the allowed radius significantly at all neutron star masses greater than $1.0 M_\odot$. See Table 4 for details. In particular, the $\pm 1\sigma$ radius range at $M = 1.4 M_\odot$, spanning all three EOS frameworks, is just 11.79–13.11 km, or $\pm 5.3\%$. In addition, although the direct measurement that we report of the radius of PSR J0740+6620 has a $\pm 1\sigma$ range of ≈ 12.2 –16.3 km, when we include other measurements the $\pm 1\sigma$ range, spanning all three EOS frameworks, narrows to 11.60–13.11 km at $M = 2.08 M_\odot$. This is dramatically improved compared with the prior $\pm 1\sigma$ range of 10.47–14.27 km over all of our EOS frameworks.

In Figure 12 we see that the maximum mass range shifts slightly to smaller values when we include our radius measurement of PSR J0740+6620; see Table 4 for details. As the table shows, the maximum mass range still has a significant dependence on the EOS framework. In a given EOS framework, there is a weak positive correlation between the radius at $M = 2.08 M_\odot$ and the maximum mass of a non-rotating star; one reflection of this is that when our analysis of the NICER and XMM-Newton data of PSR J0740+6620 is included, the maximum mass decreases and the $+1\sigma$ radius at $M = 2.08 M_\odot$ also decreases slightly in two of the three EOS frameworks. As one implication of the slightly lowered maximum mass, consider the lighter object in the gravitational wave coalescence GW190814, which has a best mass estimate of $2.59 M_\odot$ (Abbott et al. 2020b). Using our GP EOS model, without our PSR J0740+6620 measurement, there is a 12.1% probability that this object is a slowly rotating neutron star. With our PSR J0740+6620 measurement, the probability drops to 7.5%.

In Figure 13 we show the inferred 5th and 95th percentile curves for the sound speed squared as a function of number density, for our results including our PSR J0740+6620 analysis. For comparison we also show the results for our particular GP EOS framework, using only our priors. This figure shows that the sound speed c_s is not very well constrained compared to the pressure. This is expected, because $c_s^2 = dP/d\epsilon$ and thus the sound speed is related to the slope of the pressure. A comparison of the prior to the final posterior of the sound speed for the GP EOS shows that the data are not constraining beyond $\sim 4n_s$, but do tighten the allowed range significantly at lower densities. In particular, the increase in the 5th percentile of c_s^2/c^2 , compared with the prior, at densities ~ 2 – $4n_s$, is common to all three EOS frameworks.

6. Conclusions

NICER measurements of PSR J0030+0451 and PSR J0740+6620, complemented with XMM-Newton measurements of PSR J0740+6620, have reduced the uncertainty about the EOS of high-density, cold catalyzed matter. This is especially true in the density range $n \sim 1.5$ – $5n_s$. These two radius measurements




Figure 11. The 25th to 75th percentile range (blue shading) and the 5th to 95th percentile range (pink shading) of the radius as a function of mass, for non-rotating neutron stars. The faint lines show the mass–radius relations of representative individual EOS (see the text for details). At each mass only EOSs that can reach that mass contribute to the radius posterior; thus at high masses, which require hard EOSs, the radii become larger. We use the Gaussian process EOS model for the same progression of measurements as in Figure 9, where S is the symmetry energy, M refers to the high masses of three pulsars, and Λ indicates the gravitational wave measurements of tidal deformability for GW170817 and GW190425. In each panel, we only include masses below the 95th percentile of the maximum mass for non-rotating neutron stars (see Figure 12 for a closer investigation of the maximum mass). See Table 4 for details of how our radius bounds change between EOS models and when including different observations. The agreement between the methods, particularly at the $\pm 1\sigma$ level, is another indication of improving convergence between models.

Table 4
Summary of Maximum Mass and Radii at $1.4 M_{\odot}$ and $2.08 M_{\odot}$

EOS Model	Measurements	$M_{\max} (M_{\odot})$			$R_e(1.4 M_{\odot}), \text{ km}$			$R_e(2.08 M_{\odot}), \text{ km}$		
		-1σ	Median	$+1\sigma$	-1σ	Median	$+1\sigma$	-1σ	Median	$+1\sigma$
Gaussian	S, M, Λ	2.12	2.27	2.52	10.92	12.23	12.93	10.80	11.97	12.86
	+J0030	2.13	2.30	2.54	11.88	12.51	13.02	11.39	12.27	12.95
	+J0740	2.08	2.23	2.47	12.17	12.63	13.11	11.60	12.28	12.88
Spectral	S, M, Λ	2.28	2.55	2.88	10.41	11.43	12.40	10.79	11.81	12.97
	+J0030	2.39	2.78	2.93	11.52	12.22	12.67	11.56	12.77	13.11
	+J0740	2.23	2.74	2.92	11.79	12.30	12.84	11.83	12.78	13.11
Piecewise polytrope	S, M, Λ	2.13	2.37	2.64	11.26	12.32	12.89	11.04	12.21	12.90
	+J0030	2.14	2.41	2.65	11.95	12.47	12.94	11.42	12.35	12.94
	+J0740	2.09	2.27	2.61	12.16	12.56	13.01	11.67	12.36	12.91

Note. Maximum gravitational masses and equatorial circumferential radii at $M = 1.4 M_{\odot}$ and $M = 2.08 M_{\odot}$ (the best estimate of the mass of PSR J0740+6620), all at $\pm 1\sigma$ for nonrotating stars, inferred using our three EOS frameworks with three different sets of measurements. The S, M, Λ set includes constraints on the symmetry energy, the high masses of three pulsars, and the two LIGO/Virgo tidal deformability measurements. The +J0030 set also includes the Miller et al. (2019) measurement of the radius and mass of PSR J0030+0451. The +J0740 data add our measurement of the radius of PSR J0740+6620. We see that the radius estimates tighten with addition of more data, and in particular that the final $\pm 1\sigma$ radius range for a $2.08 M_{\odot}$ star, spanning all three EOS frameworks (11.60–13.11 km), is very similar to the final $\pm 1\sigma$ radius range for a $1.4 M_{\odot}$ star, spanning all three EOS frameworks (11.79–13.11 km).




Figure 12. Posterior probability distribution for the maximum gravitational mass of a non-rotating neutron star, normalized in each case so that the integral of the probability density in the $M_{\max} = 1.8\text{--}3.0 M_{\odot}$ range is 1. The highest maximum mass allowed in our analysis was $3.0 M_{\odot}$. We use the Gaussian process model for the same progression of measurements as in Figure 9, where S is the symmetry energy, M refers to the high masses of three pulsars, and Λ indicates the gravitational wave measurements of tidal deformability for GW170817 and GW190425. See Table 4 for details of how our maximum mass bounds change between EOS models and when including different observations.

also shift slightly the credible region for the maximum gravitational mass of a non-rotating neutron star.

Most importantly, our new measurement narrows significantly the inferred radius range for $1.4 M_{\odot}$ neutron stars. For instance, prior to our measurement, when we used our GP EOS model, inclusion of nuclear data, information about neutron star tidal deformability from the gravitational wave event GW170817, and our previous mass and radius measurement of PSR J0030+0451 had constrained the radius to 11.2–13.3 km at 90% credibility and 11.9–13.0 km at 68% credibility. When we also include our PSR J0740+6620 measurement, these ranges shrink to 11.8–13.4 km (90%) and 12.2–13.1 km (68%). That is, as expected, the main effect of our PSR J0740+6620 measurement is to increase the lower limit to the radius. Indeed, even when we consider the radius range spanned by all three of our EOS models, the range is 11.8–13.1 km at 68% credibility, for a fractional uncertainty of $\pm 5.3\%$. This is broadly consistent with previous radius estimates based both on nuclear theory and tidal deformabilities inferred from gravitational wave observations (Abbott et al. 2017, 2020a; De et al. 2018) and on the NICER mass and radius estimates for PSR J0030+0451 (Miller et al. 2019; Riley et al. 2019): the total radius range spanned by all the computed $\pm 1\sigma$ radius bounds, prior to our measurement of the radius of PSR J0740+6620, for $1.4 M_{\odot}$ neutron stars is $\approx 11\text{--}13.5$ km (e.g., Alvarez-Castillo et al. 2000; Dietrich et al. 2020; Greif et al. 2020; Jiang et al. 2020; Li et al. 2020, 2021; Zhao & Lattimer 2020; Essick et al. 2020b; Biswas et al. 2021; Han et al. 2021; Selva et al. 2021; Sen & Guha 2021).

Data are still being taken actively; for example, it is likely that by the end of the NICER mission the exposure times on PSR J0030+0451 and PSR J0740+6620 will roughly double




Figure 13. The 5th to 95th percentile range of the sound speed squared (in units of c^2), for all three of our models of the EOS: the red line is for the Gaussian process (GP) EOS, the blue line is for the spectral EOS, and the black line is for the piecewise polytropic EOS. The solid lines show the results when we include all measurements: symmetry energy, high-mass pulsars, tidal deformability from gravitational wave measurements, and measurements of the mass and radius of PSR J0030+0451 and PSR J0740+6620. The dotted line shows our GP results using only our priors. Because $c_s^2 = dP/d\epsilon$, the sound speed is the slope of the pressure-density relation and is therefore not as well constrained as the relation itself. We see from this plot that in our particular GP EOS framework the sound speed automatically approaches the speed of light at high densities; the prior dominates at densities $\gtrsim 4n_s$ (other choices of GP hyperparameters need not lead to this behavior; see for example Figure 2 of Landry et al. 2020). However, the rise in the 5th percentile of the sound speed, compared with the prior, at densities $\sim 2\text{--}4n_s$, is common to all three EOS frameworks.

compared with those used in current analyses, and an analysis of NICER data on PSR J0437–4715 (which has the highest X-ray flux of any non-accreting neutron star) is in progress. This gives reason for optimism that future NICER observations, and in the coming years more gravitational-wave observations of tidal deformability, will provide even clearer insight into a realm of matter which cannot be explored in laboratories.

This work was supported by NASA through the NICER mission and the Astrophysics Explorers Program. The authors acknowledge the University of Maryland supercomputing resources (<http://hpcc.umd.edu>) that were made available for conducting the research reported in this Letter. The authors also thank Felix Fuerst for facilitating the XMM-Newton director’s discretionary time observations of PSR J0740+6620, and Brad Cenko for early Swift observations. MCM thanks the Radboud Excellence Initiative for supporting his stay at Radboud University and was supported in part by NASA ADAP grant 80NSSC21K0649. MCM also thanks Gordon Baym, Cecilia Chirenti, and Sanjay Reddy for comments on an earlier version of this manuscript. W.C.G.H. appreciates use of computer facilities at the Kavli Institute for Particle Astrophysics and Cosmology and acknowledges support through grant 80NSSC20K0278 from NASA. T.T.P. is a NANOGrav Physics Frontiers Center Postdoctoral Fellow funded by the

National Science Foundation award number 1430284. I.H.S. is supported by an NSERC Discovery grant and by the Canadian Institute for Advanced Research. Support for H.T.C. was provided by NASA through the NASA Hubble Fellowship Program grant #HST-HF2-51453.001 awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., for NASA, under contract NAS5-26555. S.M.M. thanks the Natural Sciences and Engineering Research Council of Canada for funding. Portions of this work performed at NRL were funded by NASA. The authors acknowledge the use of NASA's Astrophysics Data System (ADS) Bibliographic Services and the arXiv.

Facility: NICER (Gendreau et al. 2016), XMM-Newton (Strüder et al. 2001; Turner et al. 2001), Swift (Gehrels et al. 2004).

Software: emcee (Foreman-Mackey et al. 2013), MultiNest (Feroz et al. 2009), Python and NumPy (Oliphant 2007), Matplotlib (Hunter 2007), Cython (Behnel et al. 2011), schwimmbad (Price-Whelan & Foreman-Mackey 2017), and the XMM-Newton Scientific Analysis System (SAS, Gabriel et al. 2004).

Appendix A

Posterior Distributions from Analysis of Synthetic J0740-like NICER and XMM-Newton Data

Table 5 and Figure 14 show the best fit and uncertainties for a two-circle fit to a synthetic NICER plus XMM-Newton data set which mimics the actual data on PSR J0740+6620.

Table 5
Fits to Synthetic NICER and XMM-Newton Data

Parameter	Median	-1σ	$+1\sigma$	-2σ	$+2\sigma$	Assumed value
R_e (km)	13.933	12.344	17.396	11.342	22.215	13.378
$GM/c^2 R_e$	0.221	0.177	0.248	0.139	0.267	0.230
M (M_\odot)	2.081	1.991	2.174	1.895	2.261	2.080
θ_{e1} (rad)	1.618	0.892	2.341	0.583	2.602	1.820
$\Delta\theta_1$ (rad)	0.121	0.078	0.180	0.051	0.259	0.061
$kT_{\text{eff},1}$ (keV)	0.085	0.076	0.097	0.068	0.111	0.107
θ_{e2} (rad)	1.574	0.881	2.313	0.586	2.604	2.280
$\Delta\theta_2$ (rad)	0.121	0.078	0.179	0.050	0.262	0.086
$kT_{\text{eff},2}$ (keV)	0.085	0.076	0.097	0.068	0.111	0.100
$\Delta\phi_2$ (cycles)	0.554	0.422	0.578	0.409	0.592	0.422
θ_{obs} (rad)	1.526	1.467	1.585	1.444	1.610	1.463
N_H (10^{20} cm^{-2})	3.067	1.377	4.389	0.302	4.915	0.226
d (kpc)	1.211	1.025	1.400	0.840	1.598	1.099

Note. The one-dimensional credible intervals for each of the parameters in a pulse waveform model with two possibly different uniform circular spots, obtained by jointly fitting the model to synthetic NICER data (in channels 30–123) and synthetic XMM-Newton data that together mimic the actual NICER and XMM-Newton data. The rightmost column lists the parameter values that were used to generate the synthetic data. These values were taken from a model that gives a good fit to the actual PSR J0740+6620 data. The synthetic data were then generated by Poisson draws from the fluxes predicted by that model, including the background fluxes it predicts. Consistent with statistical expectations, six of the $12 \pm 1\sigma$ credible intervals, 11 of the $12 \pm 2\sigma$ credible regions, and all (12 of 12) of the $\pm 3\sigma$ credible regions contain the values of the model parameters that were assumed in generating the synthetic data.




Figure 14. Posterior probability density distributions for a model with two uniform circular spots that is fit jointly to the synthetic NICER and XMM-Newton data. The dotted lines in the one-dimensional plots for gravitational mass and distance indicate the priors that we applied. The vertical dashed lines in the one-dimensional plots indicate the value of the associated parameter that was assumed in the construction of the synthetic data. Similarly, the star symbols in the two-dimensional plots show the assumed values of both parameters. In this and the other corner plots in the appendices, brighter colors in the two-dimensional plots indicate higher posterior probability densities.

Appendix B

Posterior Distributions from Analysis of Only NICER Data

Table 6 and Figure 15 show the best fit and uncertainties for a two-circle fit to just the NICER data on PSR J0740+6620.

Table 6
Fits to Only NICER Data

Parameter	Median	-1σ	$+1\sigma$	-2σ	$+2\sigma$	Best fit
R_e (km)	11.512	10.382	13.380	9.642	16.256	11.008
GM/c^2R_e	0.266	0.229	0.293	0.188	0.308	0.278
M (M_\odot)	2.072	1.978	2.159	1.885	2.247	2.071
θ_{e1} (rad)	1.655	0.667	2.541	0.292	2.858	0.494
$\Delta\theta_1$ (rad)	0.329	0.205	0.541	0.132	0.899	0.328
$kT_{\text{eff},1}$ (keV)	0.072	0.064	0.079	0.057	0.087	0.078
θ_{e2} (rad)	1.582	0.641	2.510	0.310	2.847	1.803
$\Delta\theta_2$ (rad)	0.326	0.204	0.539	0.132	0.854	0.208
$kT_{\text{eff},2}$ (keV)	0.072	0.064	0.079	0.057	0.087	0.079
$\Delta\phi_2$ (cycles)	0.527	0.441	0.560	0.422	0.578	0.550
θ_{obs} (rad)	1.528	1.468	1.586	1.444	1.610	1.535
N_H (10^{20} cm^{-2})	1.833	0.599	3.414	0.094	4.667	0.858
d (kpc)	1.213	1.025	1.407	0.843	1.596	1.086

Note. One-dimensional credible regions, and best fit, obtained by fitting a model with two possibly different uniform circular spots to only the NICER data, using energy channels 30–123. These credible regions may be compared with the regions in Table 7, where both NICER and XMM-Newton data are analyzed. The analysis here was performed using a free background in each NICER energy channel; as discussed in Section 4.2, future incorporation of reliable background models for NICER will likely increase the NICER-only radius to values more comparable to what we infer when we also use XMM-Newton data.




Figure 15. Posterior probability density distributions for a model with two uniform circular spots that is fit to only the NICER data. Note that because the spots do not overlap with each other, their roles can be swapped. Moreover, because the observer inclination is close to the equator, there is a near-degeneracy between spots in the northern hemisphere and spots in the southern hemisphere. The dotted lines in the one-dimensional plots for gravitational mass and distance indicate the priors that we applied.

Appendix C

Posterior Distributions from Analysis of NICER and XMM-Newton Data

Table 7 and Figure 16 show the best fit and uncertainties for a two-circle fit to the NICER plus XMM-Newton data on PSR J0740+6620.

Table 7
Fits to NICER and XMM-Newton data

Parameter	Median	-1σ	$+1\sigma$	-2σ	$+2\sigma$	Best fit
R_e (km)	13.713	12.209	16.326	11.243	20.051	13.823
GM/c^2R_e	0.222	0.187	0.249	0.151	0.267	0.221
M (M_\odot)	2.062	1.971	2.152	1.883	2.242	2.067
θ_{c1} (rad)	1.600	0.900	2.319	0.586	2.631	0.834
$\Delta\theta_1$ (rad)	0.098	0.065	0.145	0.044	0.219	0.087
$kT_{\text{eff},1}$ (keV)	0.094	0.083	0.105	0.073	0.118	0.098
θ_{c2} (rad)	1.612	0.917	2.297	0.572	2.618	1.223
$\Delta\theta_2$ (rad)	0.096	0.064	0.143	0.042	0.213	0.066
$kT_{\text{eff},2}$ (keV)	0.094	0.083	0.106	0.073	0.120	0.103
$\Delta\phi_2$ (cycles)	0.558	0.422	0.579	0.409	0.592	0.577
θ_{obs} (rad)	1.527	1.467	1.586	1.444	1.610	1.561
N_H (10^{20} cm^{-2})	1.137	0.315	2.549	0.043	4.174	0.006
d (kpc)	1.215	1.027	1.404	0.852	1.593	1.151

Note. One-dimensional credible regions, and best fit, obtained by jointly fitting a model with two possibly different uniform circular spots to channels 30–123 of the NICER data and to the XMM-Newton data on PSR J0740+6620. This analysis used the nominal calibration of NICER and XMM-Newton. These credible regions may be compared with the regions in Table 6, where only NICER data are analyzed.




Figure 16. Posterior probability density distributions for a model with two uniform circular spots that is fit to the NICER and XMM-Newton data, using the nominal calibration for NICER and XMM-Newton. The dotted lines in the one-dimensional plots for gravitational mass and distance indicate the priors that we applied.

Appendix D
Posterior Distributions from Analysis of NICER and
XMM-Newton Data with an XMM-Newton Calibration
Parameter

Table 8 and Figure 17 show the best fit and uncertainties for a two-circle fit to the NICER plus XMM-Newton data on PSR J0740+6620, in which we have added a parameterized normalization for the XMM-Newton calibration.

Table 8
 Fits to NICER and XMM-Newton Data with a Parameterized Normalization for the XMM-Newton Calibration

Parameter	Median	-1σ	$+1\sigma$	-2σ	$+2\sigma$	Best fit
R_e (km)	13.705	12.153	16.298	11.140	20.199	13.524
GM/c^2R_e	0.223	0.187	0.250	0.151	0.270	0.227
M (M_\odot)	2.064	1.973	2.154	1.882	2.244	2.082
θ_{e1} (rad)	1.635	0.895	2.324	0.575	2.634	2.399
$\Delta\theta_1$ (rad)	0.097	0.064	0.144	0.042	0.214	0.106
$kT_{\text{eff},1}$ (keV)	0.095	0.084	0.107	0.074	0.120	0.098
θ_{e2} (rad)	1.647	0.910	2.327	0.573	2.632	1.876
$\Delta\theta_2$ (rad)	0.097	0.065	0.143	0.044	0.212	0.076
$kT_{\text{eff},2}$ (keV)	0.095	0.084	0.106	0.074	0.119	0.102
$\Delta\phi_2$ (cycles)	0.551	0.422	0.578	0.409	0.591	0.571
θ_{obs} (rad)	1.525	1.467	1.586	1.444	1.610	1.536
N_H (10^{20} cm^{-2})	1.081	0.293	2.489	0.040	4.116	0.033
d (kpc)	1.209	1.023	1.403	0.845	1.599	1.206
A_{XMM}	0.969	0.919	1.043	0.902	1.090	0.917

Note. One-dimensional credible regions, and best fit, obtained by jointly fitting a model with two possibly different uniform circular spots to channels 30–123 of the NICER data and to the XMM-Newton data on PSR J0740+6620. Here we add a parameter, A_{XMM} , which is the energy-independent ratio of the XMM-Newton effective area to its nominal value. These credible regions are nearly identical to those in Table 7, where we assumed the nominal effective area.




Figure 17. Posterior probability density distributions for a model with two uniform circular spots that is fit to the NICER and XMM-Newton data, where we allow the XMM-Newton effective area to be multiplied by an energy-independent parameter A_{XMM} . The dotted lines in the one-dimensional plots for gravitational mass and distance indicate the priors that we applied. Here the prior on A_{XMM} was flat between 0.9 and 1.1; this range is roughly twice as broad as has been measured, as can be seen from slide 7 of <http://iaheec.org/wp-content/presentations/2020/NICER-CrossCal-IACHEC-Markwardt-2020b.pdf>, which shows that in the energy range relevant to our analyses, 0.3–1 keV, the flux inferred from 3C-273 agrees between NICER and XMM-Newton to $\sim 4\%$. Examination of this corner plot suggests that the preference for smaller A_{XMM} within this limited range is likely due to more spot sizes and configurations being allowed when the XMM-Newton data are dewighted.

ORCID iDs

- M. C. Miller <https://orcid.org/0000-0002-2666-728X>
- F. K. Lamb <https://orcid.org/0000-0002-3862-7402>
- A. J. Dittmann <https://orcid.org/0000-0001-6157-6722>
- S. Bogdanov <https://orcid.org/0000-0002-9870-2742>
- K. C. Gendreau <https://orcid.org/0000-0001-7115-2819>
- S. Guillot <https://orcid.org/0000-0002-6449-106X>
- W. C. G. Ho <https://orcid.org/0000-0002-6089-6836>
- J. M. Lattimer <https://orcid.org/0000-0002-5907-4552>
- M. Loewenstein <https://orcid.org/0000-0002-1661-4029>
- S. M. Morsink <https://orcid.org/0000-0003-4357-0575>
- P. S. Ray <https://orcid.org/0000-0002-5297-5278>

- M. T. Wolff <https://orcid.org/0000-0002-4013-5650>
- C. B. Markwardt <https://orcid.org/0000-0001-9803-3879>
- T. Okajima <https://orcid.org/0000-0002-6054-3432>
- I. Cognard <https://orcid.org/0000-0002-1775-9692>
- H. T. Cromartie <https://orcid.org/0000-0002-6039-692X>
- E. Fonseca <https://orcid.org/0000-0001-8384-5049>
- L. Guillemot <https://orcid.org/0000-0002-9049-8716>
- M. Kerr <https://orcid.org/0000-0002-0893-4073>
- A. Parthasarathy <https://orcid.org/0000-0002-4140-5616>
- T. T. Pennucci <https://orcid.org/0000-0001-5465-2889>
- S. Ransom <https://orcid.org/0000-0001-5799-9714>
- I. Stairs <https://orcid.org/0000-0001-9784-8670>

References

- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017, *PhRvL*, **119**, 161101
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2018, *PhRvL*, **121**, 161101
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2020a, *ApJL*, **892**, L3
- Abbott, R. & LIGO Scientific Collaboration, & Virgo Collaboration 2020b, *ApJL*, **896**, L44
- Adhikari, D., Albatineh, H., Androic, D., et al. 2021, *PhRvL*, **126**, 172502
- Agathos, M., Meidam, J., Del Pozzo, W., et al. 2015, *PhRvD*, **92**, 023012
- Alcock, C., & Illarionov, A. 1980, *ApJ*, **235**, 534
- Alvarez-Castillo, D., Ayriyan, A., Barnaföldi, G. G., Grigorian, H., & Pósfay, P. 2000, *EPJST*, **229**, 3615
- Alvarez-Castillo, D., Ayriyan, A., Benic, S., et al. 2016, *EPJA*, **52**, 69
- Annala, E., Gorda, T., Kurkela, A., & Vuorinen, A. 2018, *PhRvL*, **120**, 172703
- Antoniadis, J., Freire, P. C. C., Wex, N., et al. 2013, *Sci*, **340**, 448
- Arzoumanian, Z., Baker, P. T., Brazier, A., et al. 2018, *ApJ*, **859**, 47
- 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., Furusawa, S., Hatsuda, T., Kojo, T., & Togashi, H. 2019, *ApJ*, **885**, 42
- Behnel, S., Bradshaw, R., Citro, C., et al. 2011, *CSE*, **13**, 31
- Bhattacharya, D., & van den Heuvel, E. P. J. 1991, *PhR*, **203**, 1
- Bhattacharyya, A. 1943, *Bull. Calcutta Math. Soc.*, **35**, 99
- Bhattacharyya, S., Strohmayer, T. E., Miller, M. C., & Markwardt, C. B. 2005, *ApJ*, **619**, 483
- Bildsten, L., Salpeter, E. E., & Wasserman, I. 1992, *ApJ*, **384**, 143
- Binnington, T., & Poisson, E. 2009, *PhRvD*, **80**, 084018
- Biswas, B., Char, P., Nandi, R., & Bose, S. 2021, *PhRvD*, **103**, 103015
- Blaes, O. M., Blandford, R. D., Madau, P., & Yan, L. 1992, *ApJ*, **399**, 634
- Bogdanov, S., Dittmann, A. J., Ho, W. C. G., et al. 2021, *ApJL*, **94**, L15
- Bogdanov, S., Lamb, F. K., Mahmoodifar, S., et al. 2019, *ApJL*, **887**, L26
- Bogdanov, S., Rybicki, G. B., & Grindlay, J. E. 2007, *ApJ*, **670**, 668
- Buchner, J. 2014, *S&C*, **26**, 383
- Buchner, J. 2021, arXiv:2101.09675
- Capano, C. D., Tews, I., Brown, S. M., et al. 2020, *NatAs*, **4**, 625
- Carney, M. F., Wade, L. E., & Irwin, B. S. 2018, *PhRvD*, **98**, 063004
- Chia, H. S. 2020, arXiv:2010.07300
- Cierniak, M., & Blaschke, D. 2020, *EPJST*, **229**, 3663
- Contopoulos, I., & Spitkovsky, A. 2006, *ApJ*, **643**, 1139
- Cromartie, H. T., Fonseca, E., Ransom, S. M., et al. 2020, *NatAs*, **4**, 72
- Damour, T., & Nagar, A. 2009, *PhRvD*, **80**, 084035
- De, S., Finstad, D., Lattimer, J. M., et al. 2018, *PhRvL*, **121**, 091102
- Demorest, P. B., Pennucci, T., Ransom, S. M., Roberts, M. S. E., & Hessels, J. W. T. 2010, *Natur*, **467**, 1081
- Dietrich, T., Coughlin, M. W., Pang, P. T. H., et al. 2020, *Sci*, **370**, 1450
- Drischler, C., Furnstahl, R. J., Melendez, J. A., & Phillips, D. R. 2020, *PhRvL*, **125**, 202702
- Drischler, C., Han, S., Lattimer, J. M., et al. 2021, *PhRvC*, **103**, 045808
- Essick, R., Landry, P., & Holz, D. E. 2020a, *PhRvD*, **101**, 063007
- Essick, R., Tews, I., Landry, P., Reddy, S., & Holz, D. E. 2020b, *PhRvC*, **102**, 055803
- Feroz, F., Hobson, M. P., & Bridges, M. 2009, *MNRAS*, **398**, 1601
- Fonseca, E., Cromartie, H. T., Pennucci, T. T., et al. 2021, arXiv:2104.00880
- Fonseca, E., Pennucci, T. T., Ellis, J. A., et al. 2016, *ApJ*, **832**, 167
- Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, **125**, 306
- Fortin, M., Providência, C., Raduta, A. R., et al. 2016, *PhRvC*, **94**, 035804
- Gabriel, C., Denby, M., Fyfe, D. J., et al. 2004, in ASP Conf. Ser. 314, *Astronomical Data Analysis Software and Systems (ADASS) XIII*, ed. F. Ochsenbein, M. G. Allen, & D. Egret (San Francisco, CA: ASP), 759
- Gamba, R., Read, J. S., & Wade, L. E. 2020, *CQGra*, **37**, 025008
- Gehrels, N., Chincarini, G., Giommi, P., et al. 2004, *ApJ*, **611**, 1005
- Gendreau, K. C., Arzoumanian, Z., Adkins, P. W., et al. 2016, *Proc. SPIE*, **9905**, 99051H
- Greif, S. K., Hebel, K., Lattimer, J. M., Pethick, C. J., & Schwenk, A. 2020, *ApJ*, **901**, 155
- Greif, S. K., Raaijmakers, G., Hebel, K., Schwenk, A., & Watts, A. L. 2019, *MNRAS*, **485**, 5363
- Han, M.-Z., Jiang, J.-L., Tang, S.-P., & Fan, Y.-Z. 2021, arXiv:2103.05408
- Harding, A. K., & Muslimov, A. G. 2001, *ApJ*, **556**, 987
- Harding, A. K., & Muslimov, A. G. 2002, *ApJ*, **568**, 862
- Hebel, K., Lattimer, J. M., Pethick, C. J., & Schwenk, A. 2013, *ApJ*, **773**, 11
- Hernandez Vivanco, F., Smith, R., Thrane, E., et al. 2019, *PhRvD*, **100**, 103009
- Hinderer, T. 2008, *ApJ*, **677**, 1216
- Hinderer, T. 2009, *ApJ*, **697**, 964
- Ho, W. C. G., & Heinke, C. O. 2009, *Natur*, **462**, 71
- Ho, W. C. G., & Lai, D. 2001, *MNRAS*, **327**, 1081
- Hunter, J. D. 2007, *CSE*, **9**, 90
- Jiang, J.-L., Tang, S.-P., Wang, Y.-Z., Fan, Y.-Z., & Wei, D.-M. 2020, *ApJ*, **892**, 55
- Kiuchi, K., Kyutoku, K., Shibata, M., & Taniguchi, K. 2019, *ApJL*, **876**, L31
- Lackey, B. D., & Wade, L. 2015, *PhRvD*, **91**, 043002
- Lai, D. 2001, *RvMP*, **73**, 629
- Landry, P., & Essick, R. 2019, *PhRvD*, **99**, 084049
- Landry, P., Essick, R., & Chatziioannou, K. 2020, *PhRvD*, **101**, 123007
- Lattimer, J. M., & Prakash, M. 2007, *PhR*, **442**, 109
- Lattimer, J. M., & Prakash, M. 2016, *PhR*, **621**, 127
- Le Tiec, A., & Casals, M. 2021, *PhRvL*, **126**, 131102
- Le Tiec, A., Casals, M., & Franzin, E. 2021, *PhRvD*, **103**, 084021
- Li, A., Jiang, J.-L., Tang, S.-P., et al. 2020, arXiv:2009.12571
- Li, A., Miao, Z., Han, S., & Zhang, B. 2021, *ApJ*, **913**, 27
- Li, B.-A., & Magno, M. 2020, *PhRvC*, **102**, 045807
- Lindblom, L. 2010, *PhRvD*, **82**, 103011
- Lindblom, L. 2018, *PhRvD*, **97**, 123019
- Lo, K. H., Miller, M. C., Bhattacharyya, S., & Lamb, F. K. 2013, *ApJ*, **776**, 19
- Margalit, B., & Metzger, B. D. 2017, *ApJL*, **850**, L19
- Miller, M. C. 2013, arXiv:1312.0029
- Miller, M. C. 2016, *ApJ*, **822**, 27
- Miller, M. C., Chirenti, C., & Lamb, F. K. 2020, *ApJ*, **888**, 12
- Miller, M. C., & Lamb, F. K. 2015, *ApJ*, **808**, 31
- Miller, M. C., & Lamb, F. K. 2016, *EPJA*, **52**, 63
- Miller, M. C., Lamb, F. K., Dittmann, A. J., et al. 2019, *ApJL*, **887**, L24
- Most, E. R., Weih, L. R., Rezzolla, L., & Schaffner-Bielich, J. 2018, *PhRvL*, **120**, 261103
- Nättilä, J., Miller, M. C., Steiner, A. W., et al. 2017, *A&A*, **608**, A31
- Nelson, B. E., Ford, E. B., Buchner, J., et al. 2020, *AJ*, **159**, 73
- Oliphant, T. E. 2007, *CSE*, **9**, 10
- Oppenheimer, J. R., & Volkoff, G. M. 1939, *PhRv*, **55**, 374
- Özel, F., Psaltis, D., Güver, T., et al. 2016, *ApJ*, **820**, 28
- Pang, P. T. H., Dietrich, T., Tews, I., & Van Den Broeck, C. 2020, *PhRvR*, **2**, 033514
- Pavlov, G. G., & Zavlin, V. E. 1997, *ApJL*, **490**, L91
- Potekhin, A. Y. 2014, *PhyU*, **57**, 735
- Price-Whelan, A. M., & Foreman-Mackey, D. 2017, *JOSS*, **2**, 357
- Raaijmakers, G., Greif, S. K., Riley, T. E., et al. 2020, *ApJL*, **893**, L21
- Radic, D., Perego, A., Zappa, F., & Bernuzzi, S. 2018, *ApJL*, **852**, L29
- Randhawa, J. S., Meisel, Z., Giuliani, S. A., et al. 2019, *ApJ*, **887**, 100
- Rasmussen, C. E., & Williams, C. K. I. 2006, *Gaussian Processes for Machine Learning* (Cambridge, MA: MIT Press)
- Reed, B. T., Fattoyev, F. J., Horowitz, C. J., & Piekarewicz, J. 2021, *PhRvL*, **126**, 172503
- Remillard, R. A., Loewenstein, M., Steiner, J. F., et al. 2021, arXiv:2105.09901
- Rezzolla, L., Most, E. R., & Weih, L. R. 2018, *ApJL*, **852**, L25
- Ricci, C., Loewenstein, M., Kara, E., et al. 2021, arXiv:2102.05666
- Riley, T. E., Raaijmakers, G., & Watts, A. L. 2018, *MNRAS*, **478**, 1093
- Riley, T. E., Watts, A. L., Bogdanov, S., et al. 2019, *ApJL*, **887**, L21
- Riley, T. E., Watts, A. L., Ray, P. S., et al. 2021, arXiv:2105.06980
- Salmi, T., Suleimanov, V. F., Nättilä, J., & Poutanen, J. 2020, *A&A*, **641**, A15
- Selva, G., Roca-Maza, X., & Colò, G. 2021, *Symmetry*, **13**, 144
- Sen, D., & Guha, A. 2021, *MNRAS*, **504**, 3354
- Shibata, M., Zhou, E., Kiuchi, K., & Fujibayashi, S. 2019, *PhRvD*, **100**, 023015
- Silverman, B. W. 1986, *Density Estimation for Statistics and Data Analysis* (London: Chapman Hall/CRC)
- Sokal, A. 1997, *Monte Carlo Methods in Statistical Mechanics: Foundations and New Algorithms* (Boston, MA: Springer)
- Steiner, A. W., Lattimer, J. M., & Brown, E. F. 2010, *ApJ*, **722**, 33
- Stella, L., Priedhorsky, W., & White, N. E. 1987, *ApJL*, **312**, L17
- Strohmayer, T. E., & Brown, E. F. 2002, *ApJ*, **566**, 1045
- Strüder, L., Briel, U., Dennerl, K., et al. 2001, *A&A*, **365**, L18
- Tauris, T. M., & Savonije, G. J. 1999, *A&A*, **350**, 928
- Tauris, T. M., & van den Heuvel, E. P. J. 2014, *ApJL*, **781**, L13
- Tews, I., Margueron, J., & Reddy, S. 2018, *PhRvC*, **98**, 045804
- Tolman, R. C. 1939, *PhRv*, **55**, 364
- Tsai, Y.-S. 1974, *RvMP*, **46**, 815
- Tsang, M. B., Stone, J. R., Camera, F., et al. 2012, *PhRvC*, **86**, 015803
- Turner, M. J. L., Abbey, A., Arnaud, M., et al. 2001, *A&A*, **365**, L27
- 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

Wolff, M. T., Guillot, S., Bogdanov, S., et al. 2021, arXiv:2105.06978
Xie, W.-J., & Li, B.-A. 2020, [ApJ](#), 899, 4
Xie, W.-J., & Li, B.-A. 2021, [PhRvC](#), 103, 035802

Yue, T.-G., Chen, L.-W., Zhang, Z., & Zhou, Y. 2021, arXiv:2102.05267
Zhang, N.-B., & Li, B.-A. 2020, [ApJ](#), 893, 61
Zhao, T., & Lattimer, J. M. 2020, [PhRvD](#), 102, 023021