



A NICER View of the Massive Pulsar PSR J0740+6620 Informed by Radio Timing and XMM-Newton Spectroscopy

Thomas E. Riley¹ , Anna L. Watts¹ , Paul S. Ray² , Slavko Bogdanov³ , Sebastien Guillot^{4,5} , Sharon M. Morsink⁶ , Anna V. Bilous⁷ , Zaven Arzoumanian⁸, Devarshi Choudhury¹ , Julia S. Deneva⁹ , Keith C. Gendreau⁸ , Alice K. Harding¹⁰ , Wynn C. G. Ho¹¹ , James M. Lattimer¹² , Michael Loewenstein^{8,13}, Renee M. Ludlam^{14,28} , Craig B. Markwardt⁸ , Takashi Okajima⁸ , Chanda Prescod-Weinstein¹⁵ , Ronald A. Remillard¹⁶ , Michael T. Wolff² , Emmanuel Fonseca^{17,18,19,20} , H. Thankful Cromartie^{21,28} , Matthew Kerr² , Timothy T. Pennucci^{22,23} , Aditya Parthasarathy²⁴ , Scott Ransom²² , Ingrid Stairs²⁵ , Lucas Guillemot^{26,27} , and Ismael Cognard^{26,27}

¹ Anton Pannekoek Institute for Astronomy, University of Amsterdam, Science Park 904, 1090GE Amsterdam, The Netherlands; T.E.Riley@uva.nl

² Space Science Division, U.S. Naval Research Laboratory, Washington, DC 20375, USA

³ Columbia Astrophysics Laboratory, Columbia University, 550 West 120th Street, New York, NY 10027, USA

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

⁵ Université de Toulouse, CNES, UPS-OMP, F-31028 Toulouse, France

⁶ Department of Physics, University of Alberta, 4-183 CCIS, Edmonton, AB, T6G 2E1, Canada

⁷ ASTRON, the Netherlands Institute for Radio Astronomy, Postbus 2, 7990 AA Dwingeloo, The Netherlands

⁸ X-Ray Astrophysics Laboratory, NASA Goddard Space Flight Center, Code 662, Greenbelt, MD 20771, USA

⁹ George Mason University, Resident at the U.S. Naval Research Laboratory, Washington, DC 20375, USA

¹⁰ Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM 87545, USA

¹¹ Department of Physics and Astronomy, Haverford College, 370 Lancaster Avenue, Haverford, PA 19041, USA

¹² Department of Physics and Astronomy, Stony Brook University, Stony Brook, NY 11794-3800, USA

¹³ Department of Astronomy, University of Maryland, College Park, MD 207423, USA

¹⁴ Cahill Center for Astronomy and Astrophysics, California Institute of Technology, Pasadena, CA 91125, USA

¹⁵ Department of Physics and Astronomy, University of New Hampshire, Durham, New Hampshire 03824, USA

¹⁶ MIT Kavli Institute for Astrophysics & Space Research, MIT, 70 Vassar St., Cambridge, MA 02139, USA

¹⁷ Department of Physics, McGill University, 3600 rue University, Montréal, QC H3A 2T8, Canada

¹⁸ McGill Space Institute, McGill University, 3550 rue University, Montréal, QC H3A 2A7, Canada

¹⁹ Department of Physics and Astronomy, West Virginia University, P.O. Box 6315, Morgantown, WV 26506, USA

²⁰ Center for Gravitational Waves and Cosmology, West Virginia University, Chestnut Ridge Research Building, Morgantown, WV 26505, USA

²¹ Cornell Center for Astrophysics and Planetary Science and Department of Astronomy, Cornell University, Ithaca, NY 14853, USA

²² National Radio Astronomy Observatory, 520 Edgemont Road, Charlottesville, VA 22903, USA

²³ Institute of Physics, Eötvös Loránd University, Pázmány P.s. 1/A, 1117 Budapest, Hungary

²⁴ Max Planck Institute for Radio Astronomy, Auf dem Hügel 69, D-53121 Bonn, Germany

²⁵ Dept. of Physics and Astronomy, University of British Columbia, 6224 Agricultural Road, Vancouver, BC V6T 1Z1, Canada

²⁶ Laboratoire de Physique et Chimie de l'Environnement et de l'Espace—Université d'Orléans/CNRS, F-45071, Orléans Cedex 02, France

²⁷ Station de radioastronomie de Nançay, Observatoire de Paris. CNRS/INSU, F-18330, Nançay, France

Received 2021 April 16; revised 2021 June 9; accepted 2021 June 9; published 2021 September 8

Abstract

We report on Bayesian estimation of the radius, mass, and hot surface regions of the massive millisecond pulsar PSR J0740+6620, conditional on pulse-profile modeling of Neutron Star Interior Composition Explorer X-ray Timing Instrument event data. We condition on informative pulsar mass, distance, and orbital inclination priors derived from the joint North American Nanohertz Observatory for Gravitational Waves and Canadian Hydrogen Intensity Mapping Experiment/Pulsar wideband radio timing measurements of Fonseca et al. We use XMM-Newton European Photon Imaging Camera spectroscopic event data to inform our X-ray likelihood function. The prior support of the pulsar radius is truncated at 16 km to ensure coverage of current dense matter models. We assume conservative priors on instrument calibration uncertainty. We constrain the equatorial radius and mass of PSR J0740+6620 to be $12.39^{+1.30}_{-0.98}$ km and $2.072^{+0.067}_{-0.066} M_{\odot}$ respectively, each reported as the posterior credible interval bounded by the 16% and 84% quantiles, conditional on surface hot regions that are non-overlapping spherical caps of fully ionized hydrogen atmosphere with uniform effective temperature; a posteriori, the temperature is $\log_{10}(T [\text{K}]) = 5.99^{+0.05}_{-0.06}$ for each hot region. All software for the X-ray modeling framework is open-source and all data, model, and sample information is publicly available, including analysis notebooks and model modules in the Python language. Our marginal likelihood function of mass and equatorial radius is proportional to the marginal joint posterior density of those parameters (within the prior support) and can thus be computed from the posterior samples.

Unified Astronomy Thesaurus concepts: Millisecond pulsars (1062); Rotation powered pulsars (1408); Pulsars (1306); Radio pulsars (1353); X-ray astronomy (1810); Neutron stars (1108)

Supporting material: animation, figure sets

1. Introduction

The nature of supranuclear density matter, as found in neutron star cores, is highly uncertain. Possibilities include

²⁸ NASA Hubble Fellowship Program Einstein Postdoctoral Fellow.

both neutron-rich nucleonic matter and stable states of strange matter in the form of hyperons or deconfined quarks (for recent reviews see Oertel et al. 2017; Baym et al. 2018; Tolos & Fabbietti 2020; Yang & Piekarewicz 2020; Hebel 2021). One way to determine the dense matter equation of state (EOS, a function of both composition and inter-particle interactions) is to measure neutron star masses and radii (Lattimer & Prakash 2016; Özel & Freire 2016). There are several possible methods, but in this Letter we focus on pulse-profile modeling (see Watts et al. 2016; Watts 2019, and references therein). This requires precise phase-resolved spectroscopy, a technique that motivated the design and development of NASA’s Neutron Star Interior Composition Explorer (NICER).

The NICER X-ray Timing Instrument (XTI) is a payload installed on the International Space Station. The primary observations carried out by NICER are order megasecond exposures of rotation-powered X-ray millisecond pulsars (MSPs) that may be either isolated or in a binary system (Bogdanov et al. 2019a). Surface X-ray emission from the heated magnetic poles propagates to the NICER XTI through the curved spacetime of the pulsar, and the compactness affects the signal registered by the instrument. However, these pulsars also spin at relativistic rates. So with a precisely measured spin frequency derived from radio timing and high-quality spectral-timing event data, we are also sensitive to rotational effects on surface X-ray emission, and therefore to the radius of the pulsar independent of the compactness (see Bogdanov et al. 2019b and references therein).

The first joint mass and radius inferences conditional²⁹ on pulse-profile modeling of NICER observations of an MSP were reported by Miller et al. (2019) and Riley et al. (2019).

The target was PSR J0030+0451, an isolated³⁰ source spinning at approximately 205 Hz. Being isolated, the radio timing model for this MSP has no dependence on its mass, in contrast to the radio timing model for an MSP in a binary. This meant that a wide prior on the mass had to be assumed in the pulse-profile modeling, which nevertheless—due to the high quality of the data set in terms of the number of pulsed counts—delivered credible intervals on the mass and radius posteriors at the $\sim 10\%$ level. These posterior distributions have been used to infer properties of the dense matter EOS (in combination with constraints from radio timing, gravitational wave observations, and nuclear physics experiments). To give a few examples, there have been follow-on studies constraining both parameterized EOS models (Miller et al. 2019; Raaijmakers et al. 2019, 2020; Dietrich et al. 2020; Jiang et al. 2020; Al-Mamun et al. 2021) and non-parameterized EOS models (Essick et al. 2020; Landry et al. 2020), some focusing particularly on the neutron star maximum mass (Lim et al. 2020; Tews et al. 2021). Others have focused on specific nuclear physics questions: hybrid stars and phase transitions to quark matter (Alvarez-Castillo et al. 2020; Blaschke et al. 2020; Christian & Schaffner-Bielich 2020; Li et al. 2020; Tang et al. 2021; Xie & Li 2021), the three-nucleon potential (Maselli et al. 2021), relativistic mean-field models (Traversi et al. 2020), muon fraction content (Zhang & Li 2020a), and the nuclear symmetry energy (Zimmerman et al. 2020; Biswas et al. 2021).

This is by no means an exhaustive review of the literature, but serves to give a flavor of how the previous NICER result has been used.

Pulse-profile modeling also yields posterior distributions for the properties of the hot X-ray-emitting regions on the star’s surface, which are assumed to be related to the star’s magnetic field structure (Pavlov & Zavlin 1997). The analysis of PSR J0030+0451 implied a complex non-dipolar field (Bilous et al. 2019) and the posteriors have been used in follow-on studies of pulsar magnetospheres and radiation mechanisms (Chen et al. 2020; Suvorov & Melatos 2020; Kalapotharakos et al. 2021).

The subject of this Letter is the rotation-powered millisecond pulsar PSR J0740+6620, spinning at approximately 346 Hz as it orbits with a binary companion (Cromartie et al. 2020). Being in a binary at a favorable inclination for measurement of the Shapiro delay allows the mass of this source to be measured independently via radio timing (e.g., Lorimer & Kramer 2004). Cromartie et al. (2020) reported a mass of $2.14_{-0.09}^{+0.10} M_{\odot}$, making this the highest-mass (well-constrained) neutron star. High-mass neutron stars (with the highest central densities) are particularly powerful in terms of their potential to constrain the dense matter EOS. The mass alone can cut down parameter space, but a measurement of radius adds far more (see for example Han & Prakash 2020; Xie & Li 2020).

The North American Nanohertz Observatory for Gravitational Waves (NANOGrav) and Canadian Hydrogen Intensity Mapping Experiment (CHIME)/Pulsar collaborations recently joined forces to perform wideband radio timing of PSR J0740+6620 (Fonseca et al. 2021a). They derived an updated measurement of the pulsar mass ($2.08 \pm 0.07 M_{\odot}$), its distance from Earth, and the orbital inclination. From these informative measurements and NICER observations (Wolff et al. 2021) comes the potential for synergistic constraints on X-ray pulse-profile parameters that do not appear in the wideband radio timing solution, in particular the radius of PSR J0740+6620 and the properties of the hot surface X-ray-emitting regions conditional on a model. We report such inferences in this Letter.

We organize this Letter as follows. In Section 2 we summarize the pulse-profile model components; we provide additional detail about novel model components to augment the information in Riley et al. (2019) and Bogdanov et al. (2019b). Section 2 also covers the X-ray likelihood function—which is the probability of the NICER XTI event data and the spectroscopic event data acquired by the XMM-Newton European Photon Imaging Camera (EPIC)—and details about the prior probability density functions (PDFs) of model parameters that are fundamentally shared by all models or multiple models. In Section 3 we report model inferences and details about the prior PDFs that are specific to a given surface hot region model. In Section 4 we discuss these inferences in detail, covering their physical implications and potential systematic errors. In Section 5 we conclude by reporting the mass and radius constraints and commenting on the outlook for future pulse-profile modeling efforts. EOS inference using our derived mass–radius posterior is carried out in a companion paper (Raaijmakers et al. 2021).

2. Modeling Procedure

The methodology in this Letter is largely shared with that of Riley et al. (2019) and Bogdanov et al. (2019b, 2021). In this

²⁹ For an introduction to the concept of conditional probabilities within Bayesian inference see Sivia & Skilling (2006), Trotta (2008), Hogg (2012), Hogg et al. (2020), Gelman et al. (2013), and Clyde et al. (2021).

³⁰ No binary companion has ever been detected despite 20 years of intensive radio timing (Lommen et al. 2000; Arzoumanian et al. 2018).

section, we summarize that methodology and give a more detailed report of the new model components.

We formulate the general form of the likelihood function shared by all models and also the prior PDFs of parameters that are shared by all models. We reserve definitions of prior PDFs of phenomenological surface hot region parameters for Section 3 where posterior inferences are reported. All posterior PDFs are computed using the X-ray Pulse Simulation and Inference (X-PSI) v0.7 framework (Riley 2021), an updated version of the package used by Riley et al. (2019).³¹ The analysis files may be found in the persistent repository of Riley et al. (2021); the data products; the numeric model files including the telescope calibration products; model modules in the Python language using the X-PSI framework; posterior sample files; and Jupyter analysis notebooks. We begin by introducing the data sets used in the analysis, including the most important aspects of the data selection and preparation.

2.1. X-Ray Event Data

In this section we summarize the event data sets reported by Wolff et al. (2021), including any pre-processing tailored to this present Letter.

2.1.1. NICER XTI

The PSR J0740+6620 analysis is based on a sequence of exposures with NICER XTI (hereafter NICER) acquired in the period 2018 September 21–2020 April 17. The event data were obtained using similar filtering criteria to the previously analyzed NICER data set of PSR J0030+0451. We only used good time intervals when all 52 active detectors were collecting data but rejected all events from DetID 34, as it often shows enhanced count rates relative to the other 51 detectors. We excluded time intervals when PSR J0740+6620 was situated at an angle $\leq 80^\circ$ from the Sun to reduce the increase in background in the lowest channels due to optical loading. We further excised events collected at low cut-off rigidity (COR_SAX values < 5) to minimize particle background contamination. The resulting cleaned event list has an on-source exposure time of 1602683.761 s.

We model events registered in the PI channel subset [30, 150], corresponding to the nominal photon energy range [0.3, 1.5] keV.³² The quoted nominal photon energy for a channel is the energy mid-point for that channel.

Below nominal energy 0.3 keV, the instrument calibration products have greater a priori uncertainty due to the sharp lower-energy threshold cutoff, as well as the presence of unrejected instrumental noise that can extend above 0.25 keV. We therefore neglect the information below channel 30 in order to reduce the risk of inferential bias. Wolff et al. (2021) detect pulsations with the highest significance when considering event data up to nominal energy ~ 1.2 keV; in this Letter we include the additional information at nominal energies in the range [1.2, 1.5] keV. The number of counts generated by the PSR J0740+6620 hot regions in channels 150 and above, however, is small and diminishes relative to the counts generated by background processes (including a non-thermal component from the environment in the near-vicinity of PSR J0740+6620); we

³¹ <https://github.com/ThomasEdwardRiley/xpsi>

³² A photon that deposits all of its energy, $E \in [0.3, 1.5]$ keV, in a detector with perfect energy resolution is registered in channel subset [30, 150] after mapping according to the gain-scale calibration product.




Figure 1. Phase-folded PSR J0740+6620 event data for two rotational cycles (for clarity): we use 32 phase intervals (bins) per cycle and the count numbers in bins separated by one cycle—in a given channel—are identical. The total number of counts is given by the sum over all phase-channel pairs. The top panel displays the pulse profile summed over the contiguous subset of channels [30, 150]. We use the red bar to indicate the typical standard deviation that an adequately performing Poisson count model will exhibit. The bottom panel displays the phase-channel-resolved count numbers for channel subset [30, 150]. For likelihood function evaluation we group all event data registered in a given channel into phase intervals spanning a single rotational cycle. The description in this caption is based on that given by Riley et al. (2019) about the corresponding count-number figure for PSR J0030+0451.

therefore neglect this higher-energy information, and focus energy resolution at nominal energies below 1.5 keV.

The NICER event data are phase-folded according to the NANOGrav radio timing solution presented in Fonseca et al. (2021a). The phase-binned count numbers are displayed in Figure 1.

2.1.2. XMM-Newton EPIC

The XMM-Newton (hereafter XMM) telescope observed PSR J0740+6620 as part of a Director’s Discretionary Time program in three visits: 2019 October 26 (ObsID 0851181601), 2019 October 28 (ObsID 0851181401), and 2019 November 1 (ObsID 0851181501). The EPIC instruments (pn, MOS1, and MOS2) were employed in “Full Frame” imaging mode with the “Thin” optical blocking filters in place. Due to the insufficiently fast read-out times (73.4 ms for EPIC-pn and 2.6 s for EPIC-MOS1/2), the data do not provide useful pulse timing information, so only phase-averaged spectral information is available. The XMM data were reduced using the Science Analysis Software (SAS³³) using the standard set of analysis threads.

The event data were first screened for periods of strong soft proton background flaring. The resulting event lists were then further cleaned by applying the recommended PATTERN (≤ 12 for MOS1/2 and ≤ 4 for pn) and FLAG (0) filters. The final source event lists were obtained by extracting events from a circular region of radius $25''$ centered on the radio timing position of PSR J0740+6620. This resulted in effective

³³ 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.

exposure times of 6.81, 17.96, and 18.7 ks for the pn, MOS1, and MOS2 instruments, respectively.

2.2. Radiation Propagation from Surface to Telescope

2.2.1. Design of the Equatorial Radius Prior

One of the ultimate aims of this Letter is to report a likelihood function of mass and (equatorial) radius that can be used for EOS posterior computation. As highlighted by Riley et al. (2018), it is desirable to define a prior PDF which is jointly flat with respect to two parameters within the prior support; these parameters can simply be mass and radius, or some deterministically related variables, such as mass and compactness. The marginal joint posterior PDF of these parameters is then proportional to the marginal likelihood function of these parameters, meaning that the marginal likelihood function can be estimated from posterior samples, for use in subsequent inferential analyses.

In this Letter we follow Riley et al. (2019) by defining a joint prior PDF of mass and radius that is flat within the prior support, which is maximally inclusive with regard to theoretical EOS predictions. The prior support is zero for $R > 16$ km because we are not aware of any EOS models predicting a radius higher than this limit that would be compatible with current constraints from nuclear physics, or with the constraints posed by the gravitational wave measurement of tidal deformability for the binary neutron star merger GW170817 (see, e.g., Reed et al. 2021). A difference to Riley et al. (2019) is that we define the prior support using a higher compactness limit as discussed in Section 2.2.3 (see Table 1 for the prior PDF and support). The prior support is also subject to the condition that the effective gravity lies within a bounded range at every point of the rotating oblate surface, in order to conform to bounds on the atmosphere models we condition on (see Section 2.3.2 and Table 1).

In Section 2.2.2 we introduce information about the mass in the form of a marginal PDF whose shape approximates the marginal likelihood function of mass derived in a recent radio timing analysis; we multiply this likelihood function with the jointly flat prior PDF of mass and radius described above, thereby defining an updated prior PDF for the X-ray pulse-profile modeling in this Letter. Our prior PDF for pulse-profile modeling is therefore not jointly flat with respect to mass and radius, but all shape information is likelihood-based, satisfying the requirements of Riley et al. (2018).

2.2.2. Mass, Inclination, and Distance Priors

Radiation is transported from the oblate X-ray-emitting surface to distant static telescopes by relativistic ray-tracing, as described by Bogdanov et al. (2019b) and references therein. The gravitational mass of PSR J0740+6620, the distance of PSR J0740+6620 from Earth, and the inclination of the Earth's line of sight to the PSR J0740+6620 spin axis are all parameters of the source–receiver system that need to be specified in order to simulate an X-ray signal incident on a telescope, and can be inferred via pulsar radio timing. Note that these static model X-ray telescopes are fictitious constructs: the real telescopes are in motion relative to the pulsar (see, e.g., Riley 2019). The NICER event data pre-processing operations include the phase-folding of on-board arrival times according to a NANOGrav radio timing solution; the phase-folded events are then the events that would be registered by a telescope that is static relative to the pulsar.

The NANOGrav and CHIME/Pulsar collaborations jointly performed wideband radio timing of PSR J0740+6620 (Fonseca et al. 2021a, 2021b). A product of this radio timing work was a joint posterior PDF of the pulsar gravitational mass, Earth distance, and the inclination Earth subtends to the orbital direction³⁴ that we can condition on as an informative prior PDF for joint NICER and XMM X-ray pulse-profile modeling. The synergy—due to degeneracy breaking—between radio timing and X-ray modeling yields a higher sensitivity to the pulsar radius and the parameters of a surface X-ray emission model. We used two sets of joint posterior PDFs provided by Fonseca et al. (2021a) as prior information over the course of our analysis: one set that was numerically estimated using a weighted least-squares solver, used in this Letter for an exploratory analysis (Section 3.2), and a second that instead used a generalized least-squares (GLS) algorithm for determining timing model parameters from wideband timing data, used in this Letter for a production analysis (Section 3.3). Fonseca et al. (2021a) report on results obtained with GLS fitting of wideband timing data, though they publicly provide both sets of PDFs as they were used in this work.

Systematic error in the parameter estimates by NANOGrav and CHIME/Pulsar is dominated by sensitivity to the choice of the dispersion-measure variability (DMX) model. Posterior PDFs were computed for three DMX variants, and the systematic error implied by the posterior variation was substantially smaller than the formal posterior spread conditional on any one DMX model. We average (marginalize) the posterior PDFs over the three DMX models with a uniform weighting.³⁵

The prior PDFs we implement for every model that conditions on joint NANOGrav and CHIME/Pulsar wideband radio timing are as follows. The distance is separable from the pulsar mass and the Earth inclination to the orbital axis, but the mass and (cosine of the) inclination are correlated. For the distance, we implement the NANOGrav \times CHIME/Pulsar measurement by first marginalizing the PDFs over the DMX models and then reweighting from the flat distance prior conditioned on by NANOGrav and CHIME/Pulsar to a physical distance prior in the direction of PSR J0740+6620 following the exemplar treatment of distance information by Igoshev et al. (2016).³⁶ The physical distance prior, however, remains relatively flat in the context of the likelihood function—meaning the likelihood function is dominant in the measurement—and the modification to the original DMX-averaged PDF is entirely minor. Additional distance likelihood information derives from the Shklovskii effect which effectively truncates the PDF, putting an upper limit on the distance (Shklovskii 1970); such truncation, however, occurs well into the tail of the NANOGrav \times CHIME/Pulsar distance distribution, so it is also an unimportant detail. Finally, we approximate the NANOGrav \times CHIME/Pulsar marginal PDF of distance as a skewed and truncated Gaussian distribution (see Table 1 for details). We display the marginal distance PDF variants referenced above in Figure 2.

For the mass and the (cosine of the) orbital inclination we implement a multi-variate Gaussian distribution with the

³⁴ The pulsar spin angular momentum is assumed to be parallel to the orbital angular momentum, but we also test sensitivity to an isotropic spin-direction prior.

³⁵ Equivalent to choosing a prior mass function of the DMX models that yields posterior probability ratios of unity between those models.

³⁶ The ecliptic coordinates of PSR J0740+6620 reported by Arzoumanian et al. (2018) were transformed to Galactic coordinates using Astropy (<http://www.astropy.org>; Astropy Collaboration et al. 2013, 2018).




Figure 2. Implementation of the joint NANOGraV and CHIME/Pulsar posterior probability distribution function (PDF) of distance (Fonseca et al. 2021a) as a prior PDF in the context of joint NICER and XMM X-ray pulse-profile modeling. The red distributions are marginal posterior PDFs of the distance derived by NANOGraV and CHIME/Pulsar, conditional on a flat prior PDF; the solid red distribution is the posterior PDF marginalized over the three dispersion-measure variability (DMX) models, and each of the lightweight red distributions is a posterior PDF conditional on a particular DMX model. These posterior PDFs of distance are proportional to the marginal likelihood function of distance. The black dashed-dotted PDF is the prior distribution of Galactic pulsars in the direction of PSR J0740+6620, adopted from Igoshev et al. (2016) and renormalized to the interval $D \in [0.6, 2.0]$ kpc on which the NANOGraV \times CHIME/Pulsar PDF is supported. The black solid distribution is the posterior PDF of distance conditional on the black dashed-dotted prior PDF. The blue dashed distribution is an approximating PDF that we condition on as a prior PDF for the pulse-profile modeling in this Letter, after renormalizing to be supported on the interval $D \in [0.0, 1.7]$ kpc; refer to Table 1 for details needed to reproduce this approximating PDF.

covariance matrix of the DMX-averaged PDF. Implicit in this PDF is a prior PDF defined by the radio timing collaboration. The marginal prior PDF of the pulsar mass is *not* flat: the prior PDF of the cosine of the orbital inclination and of the mass of the white dwarf companion of PSR J0740+6620 is jointly flat, and these variables map deterministically to the pulsar mass through the binary mass function and (precise) radio timing of the system. We do not modify the joint prior PDF of the pulsar mass and the inclination despite the likelihood function of the pulsar mass (marginalized over all other radio timing parameters) being formally desirable for EOS parameter estimation (see Section 2.2.1 and Riley et al. 2018). In this case the posteriors are sufficiently dominated by the likelihood function to ignore the small structural modifications that would result from tweaking relatively diffuse priors. Moreover, changing the pulsar mass prior PDF would change the prior PDF of the companion, which may be undesirable.

Ultimately, handling the detailed structure³⁷ of these prior PDFs—i.e., beyond the location and marginal spread of the parameters—is not important because the NICER likelihood

³⁷ See Figure 2 for reference. The form of the DMX-marginalized joint mass and inclination prior PDF is the dominating factor in the posterior PDF rendered in Figure 7; for supplementary plots of the variation of the joint prior PDF of mass and inclination with DMX model, refer to the analysis notebooks released with this Letter, in which the model components are constructed.

function is not sufficiently sensitive to some combination of these radio-timing parameters and its native parameters (such as equatorial radius) for posterior inferences to change to any discernable degree. That is, the changes will be difficult to resolve from sampling noise and implementation error for instance (Higson et al. 2018, 2019), and relative to model-to-model systematic variations that, when marginalized over, further broaden the posteriors of shared parameters of interest. Also note that the distance only enters in the likelihood function in combination with the effective-area scaling parameter defined in Section 2.4 that operates on all X-ray instrument response models because of global calibration error; the distance to PSR J0740+6620 could therefore be combined with this absolute scaling parameter, further diminishing the importance of treating fine details in the prior PDF of the distance.

2.2.3. Relativistic Ray-tracing

The original version of the X-PSI package used to model PSR J0030+0451 (Riley et al. 2019) via relativistic ray-tracing assumed that no more than one ray connected the telescope to each point on the stellar surface. However, photons emitted from a star with compactness greater than $GM/Rc^2 = 0.284$ can have a deflection angle larger than π , which makes multiple images of small regions at the “back” of the star possible. This was not an issue for the analysis of PSR J0030+0451, because the 95% credible region only allowed compactness values in the range of $GM/Rc^2 \leq 0.171$. Additionally, the independent analysis by Miller et al. (2019) allowed for the possibility of multiply imaged regions on the star but still led to a credible range on compactness that is too small to allow for multiple images.

The radio observations of Shapiro delay in the PSR J0740+6620 binary system by NANOGraV and CHIME/Pulsar (Fonseca et al. 2021a) lead to the marginal 68% credible interval of $M = 2.08 \pm 0.07 M_\odot$. For reference, a small selection of EOSs that cover a realistic range of stiffness are shown in Figure 3. The EOSs shown include a set of three EOSs (HLPS soft, intermediate, and stiff) constructed by Hebeler et al. (2013) that span a range of stiffness allowed by nuclear experiments, as well as the A18+ $\delta(\nu)$ +UIX* EOS (Akmal et al. 1998) (abbreviated to APR). Each of the four curves shows the values of the equatorial compactness and mass for stars rotating at the rate of 346 Hz, computed using the code `rns` (Stergioulas & Friedman 1995). Figure 3 shows that for this set of EOSs, the large mass values ($> 2.0 M_\odot$) lead to a significant range of models that have large enough compactness (i.e., are above the horizontal line at 0.284) that multiple images of some part of the star are possible. Because the stars are oblate, the compactness at the spin poles is larger than the equatorial compactness, so the range of multiply imaged stars extends to slightly lower values of compactness; however, the change is not perceptible on this figure.

Due to the expectation that PSR J0740+6620 could be very compact, the X-PSI code was extended to allow multiple rays from any point to reach the telescope. This improvement to X-PSI as well as details of code validation are discussed in Bogdanov et al. (2021). X-PSI sums over the primary and the visible higher-order images of any regions on a star that lie in a multiply imaged region. The X-PSI package can typically detect up to the quaternary, quinary, or senary order depending on resolution settings. In practice only secondary images, and




Figure 3. Equatorial compactness (GM/Rc^2) vs. mass for stars spinning at a rate of 346 Hz. The relations for four example equation of state models (see text for description) are shown. The 95% interval for the mass prior is shaded, as is the region above compactness 0.284 where multiple images of the equator can appear. Note that for a range of compactness below 0.284, non-equatorial surface regions are multiply imaged due to oblateness. For the high masses of interest, multiple imaging is clearly relevant.

potentially the tertiary images for some configurations, might be important; but omission of secondary images can lead to large errors of $\mathcal{O}(10\%)$ in the light-curve calculation (see, e.g., the X-PSI documentation³⁸).

2.2.4. Interstellar Attenuation of X-Rays

The X-ray signal is attenuated to some degree by the intervening interstellar medium along the line of sight to the pulsar. In all models, the attenuation physics is parameterized solely by the neutral hydrogen column density N_{H} , relative to which the abundances of all other attenuating gaseous elements, dust, and grains are fixed by the state-of-the-art TBabs model (Wilms et al. 2000, updated in 2016). We implement this attenuation model as a one-dimensional lookup table with respect to energy at a fiducial column density, and then raise the attenuation factor to the power of the ratio of the column density to the fiducial value.

The column density along the line of sight to PSR J0740+6620 can be estimated using several techniques. The HEASARC neutral hydrogen map tool³⁹ yields $N_{\text{H}} \approx 3.5 \times 10^{20} \text{ cm}^{-2}$ given the most modern map (HI4PI Collaboration et al. 2016).⁴⁰

Using the relation between dispersion measure (DM) and N_{H} from He et al. (2013) together with the DM measurement reported by Cromartie et al. (2020), $N_{\text{H}} \approx 4.5 \times 10^{20} \text{ cm}^{-2}$. Finally, estimates based on 3D $E(B - V)$ extinction maps (Lallement et al. 2018), together with the relation between extinction and N_{H} from Foight et al. (2016), yield values of $N_{\text{H}} \approx 4.5 \times 10^{20} \text{ cm}^{-2}$.

Given that N_{H} could be as large as $4.5 \times 10^{20} \text{ cm}^{-2}$ based on these estimates, and given the large uncertainties in the

relations between N_{H} and other quantities such as the DM, a conservative prior of $N_{\text{H}} \sim U(0, 10^{21}) \text{ cm}^{-2}$ is warranted.

2.3. Surface Hot Regions

2.3.1. Temperature Field

The surface effective temperature field of a rotation-powered MSP is one of the primary sources (if not the primary source) of uncertainty a priori regarding the physical processes that generate the X-ray event data. The image of a neutron star cannot be resolved with any current X-ray telescope, so all our knowledge about the surface temperature field comes from models. These models heavily rely on the assumptions about magnetic field configuration, which is a crucial part in calculating heating by space-like or return current in some sub-region of the open field line footprints (Kalapotharakos et al. 2021) or anisotropic thermal conductivity of the outer star layers (e.g., De Grandis et al. 2020; Kondratyev et al. 2020). There is, therefore, an essentially unknown degree of complexity in the structure of the temperature field.

A given likelihood function is insensitive to complexity beyond some degree. We focus on a simple model of the surface temperature field: two disjoint hot regions that are simply connected spherical caps (when projected onto the unit sphere) wherein the effective temperature of the atmosphere is uniform. This model may be identified as $ST-U$ in Riley et al. (2019). The phase-folded NICER pulse profile is suggestive of two phase-separated hot regions which may or may not be disjoint.

We could begin with an even simpler model that restricts the hot regions to be related via antipodal reflection symmetry ($ST-S$ in Riley et al. 2019). Computing a posterior conditional on such a simple surface hot region model is justifiable, and it also delivers a lower-dimensional target distribution to check for egregious model implementation error (as reasoned by Riley et al. 2019). However, for the analysis reported in this Letter, we immediately explored the $ST-U$ model because antipodal reflection symmetry is physically unrealistic and the $ST-U$ model is sufficiently simple and inexpensive to condition on; nevertheless, if one is interested in the question of whether we are sensitive to deviations from this symmetry, we open-source the entire analysis package for this Letter, and the online X-PSI documentation offers guidance on how to condition on the $ST-S$ model.

Nodes can be added to the model space by incrementing the complexity of the surface hot regions, following Riley et al. (2019), wherein symmetries between hot regions are broken⁴¹ and temperature components are added.⁴² In this Letter we introduce a new binary flag for the atmosphere composition (hydrogen or helium as discussed in Section 2.3.2) within the (single-temperature) hot regions. However, we consider only one higher-complexity model than $ST-U$. We opt to do this because subject to limited (computational) resources, $ST-U$ performs well and is ultimately determined to be sufficiently complex.

³⁸ https://thomasedwardriley.github.io/xpsi/multiple_imaging.html

³⁹ <https://heasarc.gsfc.nasa.gov/cgi-bin/Tools/w3nh/w3nh.pl>

⁴⁰ Neutral hydrogen maps provide an integrated estimate along the line of sight through the Milky Way, and can therefore be interpreted as an approximate upper limit.

⁴¹ That is, the breaking of antipodal reflection symmetry and of shape symmetry.
⁴² All models in Riley et al. (2019), two of which we condition on in this Letter, can be labeled by the number of disjoint hot regions they define. Note that in practice the prior PDFs of temperature and solid angle subtended at the stellar center are diffuse and permit the contribution of a hot region to be negligible a posteriori in the NICER and XMM wave bands.

We now define prior PDFs that are shared by all models. Every hot region has one effective temperature component. The prior PDF of that effective temperature is flat in its logarithm, diffuse,⁴³ and is separable from the joint prior PDF of all other model parameters. Every hot region also has a colatitude coordinate and an azimuth coordinate (i.e., phase shift) at the center of a constituent spherical cap with a finite effective temperature component (Riley et al. 2019). The prior PDF of these coordinates is non-trivial because it is not separable from the prior PDF of other parameters when there are two surface hot regions in the model, as we proceed to explain.

In order to eliminate a trivial form of hot region exchange-degeneracy,⁴⁴ we define the prior support by imposing that one hot region object in the X-PSI model has a center colatitude that is always less than or equal to that of the companion hot region. We also impose that the hot regions are disjoint, meaning that they cannot overlap in the prior support, because otherwise the model cannot always be characterized by the number of disjoint hot regions; the non-overlapping condition is a function of the center colatitudes, the center azimuths, and the hot region angular radii, meaning that the joint prior is not separable, and the marginal prior PDFs are modulated relative to the PDFs that were used to construct the form of the joint PDF in six dimensions. For instance, the joint prior PDF of the azimuthal coordinates exhibits rarefaction where the azimuths are close in value.

We now construct the prior PDF specifically for the ST-U model (Section 3.3) that is the focus of this Letter. First, as a construction tool, define a flat PDF of the cosine of each hot region center colatitude, with support being the interval $\cos(\Theta) \in [-1, 1]$; such a PDF is founded on the argument of isotropy, a priori, of the direction to Earth (see below). Second, define the PDF of each hot region center azimuth⁴⁵ to be flat on the interval $2\pi\phi \in [0, 2\pi]$ radians. Third, define a flat PDF of each angular radius on the interval $[0, \pi/2]$. To calculate the marginal prior PDF of one of these parameters, marginalize over the other five parameters subject to the prior support condition of non-overlapping hot regions. For instance, the prior PDF of a hot region angular radius is given by marginalizing over the angular radius of the companion hot region, and the hot region center colatitudes and azimuths. Marginalization yields a marginal prior PDF that deviates in form from the initially flat PDF of the parameter; the support of the PDF remains unchanged, however.

Our hot region models are phenomenological and, to a degree, agnostic, despite being influenced by temperature fields implied by pulsar magnetospheric simulations. Our choice here is contrary to Riley et al. (2019), wherein the prior PDF of the colatitude at the center of a hot region in isolation—meaning one hot region on the surface—was flat. A flat PDF of colatitude, combined with a flat prior in azimuth, yields an anisotropic probability density field on the unit sphere, weighted toward polar regions. One reason this might be a sub-optimal assumption is because of X-ray selection effects: pulsed emission from two hot regions will exhibit a larger amplitude if the hot regions are closer to the equatorial zone than a polar zone, for equatorial observers as is the case assumed for

PSR J0740+6620 based on the NANOGrav \times CHIME/Pulsar wideband radio timing measurements (Section 2.2.2). However, such arguments are tenuous, taking no heed of radio selection effects and magnetospheric physics, for instance.

2.3.2. Atmosphere

The specific intensity is determined from lookup tables generated using the NSX atmosphere code assuming either fully ionized hydrogen or helium (Ho & Lai 2001) or partially ionized hydrogen (Ho & Heinke 2009).⁴⁶ In this work, we consider only fully ionized models since the limited parameter ranges of existing opacity tables (Iglesias & Rogers 1996; Badnell et al. 2005; Colgan et al. 2016) reduce the accuracy of atmosphere models employing these tables (see Bogdanov et al. 2021 and Miller et al. 2021 for further discussion and comparisons).

2.3.3. Exterior of the Hot Regions

The surface exterior to the hot regions does not explicitly radiate in any model we condition on. The signal that would be generated by a cooler exterior surface can be robustly subsumed in the phase-invariant count rate terms described in Section 2.5.1, provided that the angular scale of the hot regions (whose images are explicitly integrated over) is small and that exterior emission is soft and thus dominated by other NICER backgrounds in the channels with low nominal photon energies. Explicitly including radiation from the exterior surface would add a very weak informative mode of dependence on the mass, radius, and other parameters of interest, and thus the likelihood function is assumed insensitive to its exact treatment. Explicit treatment would also require handling of the ionization state of the cooler atmosphere.

2.4. Instrument Response Models

For NICER and each of three XMM EPIC cameras we use a tailored ancillary response file (ARF) and redistribution matrix file (RMF) to compose an on-axis response matrix. For each of the XMM cameras, the ARF and RMF are those specific to the PSR J0740+6620 extraction regions.

For NICER, we use the most recent calibration products made available by the instrument team, namely the nixtir-ef20170601v002 RMF and the nixtiaveonaxis20170601v004 ARF. For the latter, the effective areas per energy channel were rescaled by a factor of 51/52 to account for the removal of all events from DetID 34. The instrument response products (RMF and ARF files) for the XMM pn, MOS1, and MOS2 cameras tailored to the PSR J0740+6620 observations were produced with the `rmfgen` and `arfgen` commands in SAS and products from the Current Calibration Files repository.

Calibration of the performance of NICER is conducted mainly with observations of the Crab pulsar and nebula. The energy-dependent residuals in the fits to the Crab spectrum are generally $\lesssim 2\%$.⁴⁷ The calibration accuracy for the XMM MOS and pn instruments is reported to be less than 3% and 2% (at 1σ), respectively.⁴⁸ However, in the absence of a suitable absolute calibration source, the absolute energy-independent

⁴³ For example, with prior support bounds of $10^{5.1}$ K and $10^{6.8}$ K for the NSX fully ionized hydrogen atmosphere; see Section 2.3.2.

⁴⁴ Where precisely the same hot region configuration exists twice in parameter space, only one of which the prior support should include.

⁴⁵ A periodic parameter, also referred to as cyclic or wrapped.

⁴⁶ Note that only the composition within the hot regions is explicitly defined for the purpose of signal generation.

⁴⁷ See <https://heasarc.gsfc.nasa.gov/docs/heasarc/caldb/nicer/docs/xi/NICER-xi20200722-Release-Notesb.pdf> for further details.

⁴⁸ See in particular Table 1 in <https://xmmweb.esac.esa.int/docs/documents/CAL-TN-0018.pdf>.

effective area of NICER and the three XMM detectors is uncertain at an estimated level of $\pm 10\%$.

We define a free parameter that operates as an energy-independent scaling factor shared by all instruments due to the lack of a perfect astrophysical calibration source. Further, for each of NICER and XMM, we define a free parameter that also operates as an energy-independent effective area scaling. The *overall* scaling factor for each telescope is a coefficient of the response matrix, defined as the product of the shared scaling factor and the a priori statistically independent telescope-specific scaling factor. The overall scaling factors of different instruments are therefore correlated a priori to simulate—in an albeit simple way—the absolute uncertainty of X-ray flux calibration and the instrument-to-instrument calibration product errors. However, the correlated prior PDF is such that the effective area for NICER is permitted to decrease from the nominal effective area while the effective area for XMM can increase from its respective nominal effective area, and vice versa. We choose the statistically independent telescope-specific scaling factors to have equal spread a priori to the scaling factor shared by all instruments of $\pm 10\%$, therefore yielding a more conservative joint prior PDF of the overall scaling factors that we approximate as a bivariate Gaussian (see Table 1). We discuss posterior sensitivity to effective area prior information in Section 4.2.

The response models are implicitly assumed to accurately represent the time-averaged operation of the instruments during the (composite) exposures to PSR J0740+6620.

2.5. Likelihood Function and Backgrounds

In this section we formulate the likelihood function as the conditional probability of the NICER and XMM event data sets.⁴⁹ The relative constraining power offered by the XMM likelihood function is weak, but we provide detail about the methodology with the overarching aim of supporting the use of imaging observations in ongoing and future NICER pulse-profile modeling efforts.

The expected photon specific flux signal generated by the surface hot regions is calculated as a function of rotational phase (over a single rotational cycle) and energy, by integrating over the photon-specific intensity image on the sky of a distant static instrument (Bogdanov et al. 2019b). We then assume that an adequately performing model of both the NICER and XMM event data sets can be constructed by operating on the same incident signal. The count number statistics for both NICER and all XMM cameras is Poissonian: for every detector channel of the four instruments—and then for NICER every phase interval associated with a channel—the sampling distribution from which the registered count-number variate is drawn is assumed to be a Poisson distribution with an expectation that is a function of parameters of the incident signal and parameters of the model instrument response.

The NICER event time-tagging resolution is state-of-the-art, and PSR J0740+6620 is sufficiently faint that, in the absence of backgrounds, the event arrival process statistics would deviate in an entirely minor way from the incident photon

Poisson point process; in reality we contend with background radiation, and deadtime corrections to the integrated exposure time are calculated during event data pre-processing. Pile-up, the registering of multiple photons as a single event during a detector readout interval, is not an issue for XMM because the source is so faint.

As we expound upon in Sections 2.5.1 and 2.5.2, the background event data for all four instruments—after implementation of filters during data pre-processing (Wolff et al. 2021)—is modeled with a set of count rate variables, one per detector channel. As a corollary of the assumptions stated above, the background processes contributing to these events are also assumed to be Poissonian event arrival processes. The expectations of the sampling distributions of the registered count-number variates are simply sums of the expected count numbers from the PSR J0740+6620 surface hot regions and the expected background count numbers. Refer to Riley et al. (2019) for details about the NICER count-number sampling distribution; the form of the XMM camera count-number sampling distribution follows by a phase-averaging operation.

2.5.1. NICER

Let d_N be the NICER count matrix over phase interval–channel pairs. We now define variables that the NICER likelihood is a function of: let s be a vector of parameters, collected from Section 2.2, of which the incident signal generated by the pulsar is a function; let NICER denote a (parameterized) model for the response of the instrument in response to incident radiation; and let $\{\mathbb{E}[b_N]\}$ be a set of statistically independent nuisance variables, one per detector channel that contains event data to be modeled, defined as the phase-invariant expected count rate. The NICER event data are phase-resolved, so there are multiple random variates—distributed in phase—per background random variable. The background model with variables $\{\mathbb{E}[b_N]\}$ is free-form as discussed by Riley et al. (2019). The set of these variables is designed to handle complex channel-to-channel variations in the expected phase-invariant count rate. A physical background model should need (far) fewer random variables for the underlying background-generating process to capture these complexities. The likelihood function is denoted by $p(d_N|s, \{\mathbb{E}[b_N]\}, \text{NICER})$.

We numerically marginalize over the variables $\{\mathbb{E}[b_N]\}$ to yield a marginal likelihood function that is combined with a joint prior PDF to define a target distribution to draw samples from. The count rate variables have separable flat prior PDFs that are strictly improper because we do not explicitly define upper bounds on the prior support of each variable (Riley et al. 2019); the posterior is considered integrable, however, so these ill-defined prior PDFs do not result in posterior pathologies. The separable prior PDF of the count rate variables is overly diffuse, with extremely high prior-predictive complexity, such that inferences will be insensitive to minor changes to the function: $\mathbb{E}[b_N] \sim U(0, \mathcal{U})$, where the upper bound \mathcal{U} of the prior support is left unspecified.⁵⁰

The marginalization operation is separable over channels, yielding a product of one-dimensional integrals. We need to perform fast numerical marginalization over the $\{\mathbb{E}[b_N]\}$ in

⁴⁹ The domain of the likelihood function is a subset of the model parameter space (which is in general discrete–continuous mixed). The domain of the more general parameterized sampling distribution is a subset of the Cartesian product of the model parameter space and the data space; the likelihood function is the function that this sampling distribution reduces to when the data-space variables become fixed by observation.

⁵⁰ The posterior should be integrable—without proof here—even if the prior is improper, but if an upper bound were to be required, it could for instance be based on NICER count-rate limitations.

order to compress the dimensionality of the sampling space. The simpler the form of the integrands—the functions of the $\{\mathbb{E}[b_N]\}$ —the more straightforward fast numerical marginalization is. If the prior PDF $p(\mathbb{E}[b_N])$ is flat, the integrand has a single global maximum, and would be highly Gaussian if the peak of the conditional likelihood function $p(d_N|s, \mathbb{E}[b_N], \text{NICER})$ lies within the support of $p(\mathbb{E}[b_N])$.

As discussed by Riley et al. (2019), a major open question regards the total expected spectral signal that is attributable to the surface hot regions in reality.⁵¹ In lieu of a physical background model—which as discussed above is difficult to formulate for NICER—independent information conditional on data acquired with an imaging X-ray telescope such as XMM can be fundamentally valuable for deriving robust inferences if the exposure time is sufficiently long; it will, however, be substantially shorter than the order megasecond exposures of the near-dedicated NICER telescope.

Note that in order to infer the contribution from contaminating sources to the NICER event data,⁵² we would need to separate out the $\{\mathbb{E}[b_N]\}$ into an environmental background model—with some informative prior PDF including space weather contributions—and some model for the contribution from contaminating (point) sources in the field that cannot be resolved from the point-spread function (PSF) of PSR J0740+6620. However, for the principal purpose of constraining the physical properties of the pulsar, we are uninterested in the distinction between NICER backgrounds. Moreover, we do not have the statistical power to distinguish the contributions if the priors for the components are all rather diffuse. In other words, if some of the priors are informative, we do not have the statistical power to gain much information a posteriori.

2.5.2. XMM-Newton

Information about the total background contribution to the NICER event data—including contaminating (point) sources in the field (see Bogdanov et al. 2019a for a breakdown of components)—is encoded in the XMM spectroscopic imaging event data. Due to the relatively brief exposure times of the observations, very few background counts are available for a reliable background estimate. Thus, we obtained representative background estimates with higher photon statistics from the blank-sky event files provided by the XMM-Newton Science Operations Centre.⁵³ The blank-sky images were filtered in the same manner as the PSR J0740+6620 field images and the background was extracted from the same location on the detector as the pulsar. The resulting background spectrum was then rescaled so that the exposure times and BACKSCAL factors matched those of the PSR J0740+6620 exposures.

Leveraging spatial resolving power, we can derive posterior inferences about the signal from the pulsar in isolation with little confusion about the expected contributions from the pulsar and background sources.⁵⁴ When this information about the pulsar signal is injected into modeling of NICER observations—either explicitly as a prior or as a likelihood factor—the contribution from the pulsar to the NICER event

data is informed; here we work toward a likelihood function factor for the XMM telescope.

Let d_X be the XMM time-integrated count vector (over detector channels) from a region \mathcal{S} of a CCD of a camera, that is some sufficiently large subset of the support of the PSF of the PSR J0740+6620 while optimizing signal-to-noise.⁵⁵ We now define variables that the XMM likelihood is a function of. The vector of parameters of the incident signal generated by the pulsar remains as s , shared with the NICER likelihood function. Let XMM denote a (parameterized) model for the response of an XMM camera in response to incident radiation. Once more, in lieu of a physical model for the underlying background-generating process, let us define one statistically independent random nuisance variable per detector channel: the expected count rate.⁵⁶ These variables are collectively denoted $\{\mathbb{E}[b_X]\}$. The likelihood function for each XMM camera is denoted by $p(d_X|s, \{\mathbb{E}[b_X]\}, \text{XMM})$. We numerically marginalize over these model variables in the same vein as for the NICER background-marginalized likelihood function; however, to constrain the signal from PSR J0740+6620 we are in need of a more informative prior PDF of $\{\mathbb{E}[b_X]\}$.

To form the prior PDF of $\{\mathbb{E}[b_X]\}$ for an XMM camera, we consider a set of independent Poisson-random variates $\{\mathcal{B}_X\}$ defined as time-integrated astrophysical (sky) background count numbers in detector channels. The events constituting $\{\mathcal{B}_X\}$ are extracted from a blank-sky⁵⁷ region \mathcal{B} of the camera CCD array that is disjoint from region \mathcal{S} associated with PSR J0740+6620. The XMM cameras are composed of multiple side-by-side CCD detectors. For each camera, it is preferable to choose the region \mathcal{B} on the same detector as \mathcal{S} because the different CCD detectors have slightly different responses to incident radiation and therefore exhibit slightly different astrophysical (sky) backgrounds. For blank-sky exposures, the conditional sampling distribution in the space of the data is $p(\{\mathcal{B}_X\}|\{\mathbb{E}[\mathcal{B}_X]\})$, where the variables $\{\mathbb{E}[\mathcal{B}_X]\}$ are per-channel expected count numbers.

To constrain the background in the XMM images of PSR J0740+6620, we use blank-sky estimates generated using the XMM-Newton SAS tools. The blank-sky exposures are much longer in duration than the exposures to PSR J0740+6620 by almost two orders of magnitude, allowing us to constrain the astrophysical (sky) background count rate in the on-source exposures more tightly (in the absence of systematic error) than possible from blank-sky extraction regions from the same CCD during the shorter on-source exposure. For each XMM camera the respective region \mathcal{B} is not only localized to the same CCD as the PSR J0740+6620 PSF, but *is* the same region of the CCD. However, the blank-sky estimates are based on an ensemble of pointings over the sky. Our cross-check of the sky background in these reference exposures against the sky background in the vicinity of PSR J0740+6620 did not yield evidence of systematic difference.

We now formally derive the constraints on the expected number of background counts per detector channel registered within the source region \mathcal{S} of an XMM camera, $\{\mathcal{B}_X\}$, conditional on the counts registered in region \mathcal{B} over the

⁵¹ Assuming that the pulsed emission is dominated by rotationally modulated emission from the surface.

⁵² That is, contamination that cannot be robustly filtered out during event data pre-processing.

⁵³ See <https://www.cosmos.esa.int/web/xmm-newton/blank-sky>.

⁵⁴ There remains confusion, however, about what contribution from the pulsar and its near-vicinity is generated by surface hot regions.

⁵⁵ Imaging observations are configured such that the target point source is confined to a single CCD, for instance to avoid masking part of the source PSF with gaps between CCDs.

⁵⁶ Where there is now one variate per channel because the count numbers are phase-averaged.

⁵⁷ That is, a region of sky devoid of any (bright) non-diffuse X-ray sources.

ensemble of blank-sky exposures. A prior PDF of the variables $\{\mathbb{E}[\mathcal{B}_X]\}$ is needed. In the same vein as for the NICER background variables, we define a separable prior PDF

$$\mathbb{E}[\mathcal{B}_X] \sim U(0, \mathcal{V}), \quad (1)$$

where the upper-bound \mathcal{V} of the prior support is left unspecified. This trivial model over-fits the background data, with posterior PDF

$$p(\{\mathbb{E}[\mathcal{B}_X]\}|\{\mathcal{B}_X\}) \propto p(\{\mathcal{B}_X\}|\{\mathbb{E}[\mathcal{B}_X]\}); \quad (2)$$

the vector of random variates $\{\mathcal{B}_X\}$ is coincident in value with the maximum a posteriori vector in the space of $\{\mathbb{E}[\mathcal{B}_X]\}$.

As is the case for the NICER background marginalization operation described in Section 2.5.1, if the PDF $p(\{\mathbb{E}[\mathcal{B}_X]\}|\{\mathcal{B}_X\})$ is flat, the marginalization is more straightforward to execute. For example, we could form a flat prior PDF spanning the highest-density $x\%$ posterior credible interval in $\mathbb{E}[b_N]$, given $\{\mathcal{B}_X\}$. We form a flat prior PDF in a simpler way. The PDF $p(\{\mathbb{E}[\mathcal{B}_X]\}|\{\mathcal{B}_X\})$ has the approximate structure of a truncated Gaussian with deviation dependent on the number of counts \mathcal{B}_X . We let the lower and upper bounds of the support respectively be $\mathcal{L} := \max(0, \mathcal{B}_X - n\sqrt{\mathcal{B}_X})$ and $\mathcal{U} := \mathcal{B}_X + n\sqrt{\mathcal{B}_X}$, where n is a setting that controls the degree of conservatism. For some channels the number of counts is low and the PDF $p(\{\mathbb{E}[\mathcal{B}_X]\}|\{\mathcal{B}_X\})$ deviates substantially from being a truncated Gaussian, but we nevertheless adopt the same procedure—the lower-bound is pushed to zero or far into the lower tail of the distribution, but the upper bound remains sensible. For some channels, however, $\mathcal{B}_X = 0$. In these cases, we simply set $\mathcal{L} := 0$ and then iterate upwards in channel number to locate the next finite value of $\mathcal{B}_X + n\sqrt{\mathcal{B}_X}$.⁵⁸ We choose $n = 4$. With the flat PDF of $\mathbb{E}[\mathcal{B}_X]$ defined, we transform variables to derive the prior PDF $p(\mathbb{E}[b_X]|\{\mathcal{B}_X\})$ according to

$$\mathbb{E}[b_X] = \frac{A_S \mathbb{E}[\mathcal{B}_X]}{A_B T_B}, \quad (3)$$

where T_B is the blank-sky exposure time needed to transform to a count rate, and A_S/A_B is the ratio of the areas of the extraction regions \mathcal{S} (encompassing the PSR J0740+6620 PSF) and region \mathcal{B} .

In many respects this flat PDF of $\mathbb{E}[\mathcal{B}_X]$ is a remarkably conservative choice for the prior constraint on $p(\mathbb{E}[b_X]|\{\mathcal{B}_X\})$. A reason it is not conservative is the assumption of zero prior mass above an upper count-rate limit \mathcal{U} somewhere in the upper tail of the true probability PDF $p(\{\mathbb{E}[\mathcal{B}_X]\}|\{\mathcal{B}_X\})$, especially because the count-number data d_X for each XMM camera are moderately consistent with being generated by background processes. The upper limit \mathcal{U} can be decreased so that $p(\mathbb{E}[b_X]|\{\mathcal{B}_X\})$ is more informative at the risk of bias. On the other hand, \mathcal{U} can be increased, which naturally weakens the constraining power but can be justified as a safety precaution to capture a contribution such as a power-law component originating from the magnetosphere of PSR J0740+6620 that is not explicitly modeled.⁵⁹ Alternatively, we could justify a higher upper limit in terms of systematic error due to

⁵⁸ This procedure is sufficient for the channel cuts we make because there is always a higher channel with a finite number of counts \mathcal{B}_X .

⁵⁹ Note that the binary companion of PSR J0740+6620 has been inferred to be an ultra-cool white dwarf (Beronya et al. 2019) whose thermal surface emission would not contribute X-ray events, unless there is some interaction with higher-energy winds from PSR J0740+6620.

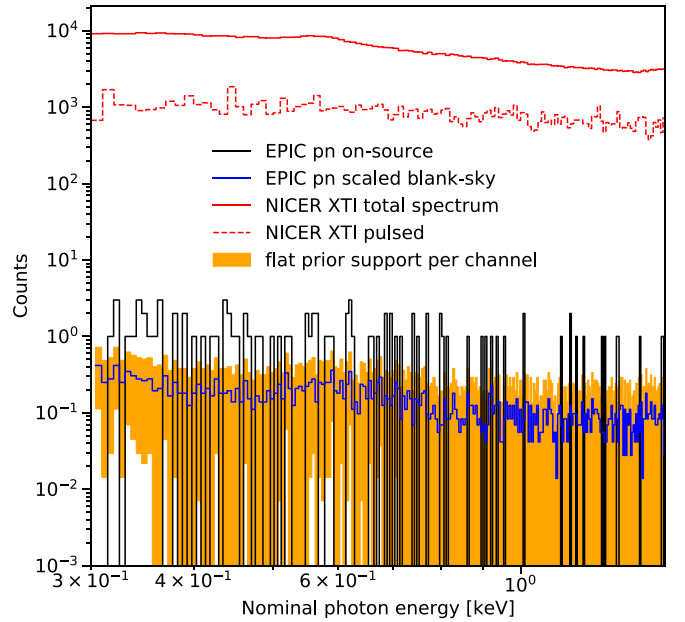


Figure 4. XMM pn camera PSR J0740+6620 count number spectrum as a function of detector channel nominal photon energy (black step function). The XMM events contain source events and events from diffuse background that can be estimated from blank-sky information. The NICER count-number spectrum is the red solid step function, including PSR J0740+6620 events and all backgrounds; the red dashed step function is the empirical pulsed count number in each channel. The XMM pn events are clearly sparse in comparison to the NICER event data, but note that the number of pulsed NICER counts is lower than indicated here, as shown in Figure 1. The blue step function is the count-number spectrum derived from blank-sky exposure, scaled down to the exposure time on PSR J0740+6620 and also scaled for the CCD extraction region area ratio A_S/A_B as described in the main text; see Equation (3). The orange shaded region is the support of the joint prior PDF of the expected count-rate variables $\{\mathbb{E}[b_X]\}$, derived from the blank-sky exposures. The figure elements rendered here are representative of the corresponding information from the XMM MOS1 and MOS2 cameras, as can be seen in the online figure set associated with Figure 16.

blank-sky estimates derived from an ensemble of exposures over the sky and variation in time of the XMM camera response matrices between the PSR J0740+6620 exposure and blank-sky exposure ensemble. The setting of $n = 4$ seems like a reasonable—albeit arbitrary—choice to balance information loss versus bias; we probe posterior sensitivity to this hyperparameter and conclude our inferences are insensitive to its value (see Section 3). We display the XMM pn camera count number spectrum in Figure 4 together with the blank-sky-derived prior PDF of the expected count-rate variables $\{\mathbb{E}[\mathcal{B}_X]\}$; we also display the NICER count-number spectrum for data quality comparison.

2.5.3. The Likelihood Function

The likelihood function given the NICER and XMM data sets may be written as

$$\begin{aligned} & p(d_N, d_X, \{\mathcal{B}_X\} | s, \{\mathbb{E}[b_N]\}, \{\mathbb{E}[b_X]\}) \\ &= p(d_N | s, \{\mathbb{E}[b_N]\}, \text{NICER}) p(d_X | s, \{\mathbb{E}[b_X]\}, \text{XMM}) \\ & \quad \times p(\{\mathcal{B}_X\} | \{\mathbb{E}[b_X]\}), \end{aligned} \quad (4)$$

where we combine the expected background count rate variables over the XMM cameras into $\{\mathbb{E}[b_X]\}$, we combine the count numbers in the PSR J0740+6620 exposures into d_X , and we combine the blank-sky count numbers into $\{\mathcal{B}_X\}$.

Introducing the flat prior densities from Equation (1), approximating the posterior PDF $p(\{\mathbb{E}[\mathcal{B}_X]\}|\{\mathcal{B}_X\})$ as flat and bounded as described in Section 2.5.2, and marginalizing over all expected background count rate variables yields the background-marginalized likelihood function

$$p(d_N, d_X, \{\mathcal{B}_X\}|s) \propto \int_{\{0\}}^{\{U\}} p(d_N|s, \{\mathbb{E}[b_N]\}, \text{NICER}) \times d\{\mathbb{E}[b_N]\} \int_{\{\mathcal{L}\}}^{\{U\}} p(d_X|s, \{\mathbb{E}[b_X]\}, \text{XMM}) d\{\mathbb{E}[\mathcal{B}_X]\}. \quad (5)$$

This function is fed as a callback to a sampling process, together with a joint prior PDF callback for the pulsar signal parameters s and parameters associated with the NICER and XMM instrument response models.

2.6. Model Space Summary

All nodes in the model space share some underlying physics. Namely, the machinery for relativistic ray-tracing: an oblate surface is embedded in an ambient Schwarzschild spacetime, and the X-ray emission emergent from the atmosphere is attenuated by the interstellar medium as it is transported to a distant static telescope. The nodes in the model space differ first and foremost in their surface hot region parameterization complexities and atmosphere composition flag values, but also in terms of the prior PDF and the likelihood function factors. The prior PDFs for the parameters controlling the shared processes (i.e., mass, equatorial radius, viewing angle to the spin axis, distance, column density) are either informative—such as the joint NANOGrav and CHIME/Pulsar measurement of mass, distance, and viewing angle—or diffuse if there is limited prior knowledge or we aim to probe the consistency of likelihood function factors across telescopes.

2.7. Posterior Computation

We implement nested sampling to compute the posterior distribution conditional on each model. Namely, we use X-PSI to construct the likelihood function and the prior PDFs and then couple them to MULTINEST (Feroz & Hobson 2008; Feroz et al. 2009; Buchner et al. 2014). For the headline model we report in this Letter, including NICER and XMM likelihood factors, the dimensionality of the sampling space is 15. Details about our nested sampling protocol are given in Riley et al. (2019; see the appendix matter in particular) and also in Riley (2019; see chapter 3 and the associated appendix). In summary, our minimum resolution settings are as follows: 10^3 live points; a bounding hypervolume expansion factor of 0.1^{-1} ; and an estimated remaining log-evidence of 10^{-1} . Regarding live points, most posteriors for sensitivity analyses are computed with 2×10^3 or 4×10^3 live points, and production calculations used 4×10^4 live points. The number of live points is the most fundamental parameter that should be changed to probe sampling resolution sensitivity—the expansion factor can be fixed at some value similar to those recommended in the literature (Feroz et al. 2009). Likelihood function evaluation time is the dominant sink, being several seconds per core for the processor speeds typical on a cluster or supercomputer. We do *not* use the constant efficiency algorithm variant for any sampling process, and we do *not* use the mode-separation

algorithm variant unless stated otherwise.⁶⁰ Regarding the mode-separation variant, if there are multiple modes of commensurate posterior mass, sampling resolution gets distributed between those modes. It follows that the bounding approximation to the constant likelihood hypersurfaces in each mode is lower than if the global resolution settings were consumed solely by that mode. We eliminate hot region exchange degeneracy from the prior support as discussed in Section 2.3.1, which eliminates mirrored modes and thus improves the resolution of that mode in terms of bounding approximation error.

For most posteriors reported in this work, the nested samples are considered high-resolution in the context of literature recommendations and, for the production analysis, were costly to generate given resource limitations. It is important to remark that our posterior computation procedure has not been validated in any meaningful way via simulation-based calibration because at present it is basically intractable for any one group to calibrate credible region coverage on a model-by-model basis (see the discussion in chapter 3 of Riley 2019 and in Riley et al. 2019). For discussion on the level of calibration we have attained by cross-checking against independent calculations performed by another group (Miller et al. 2019), we refer to Bogdanov et al. (2021). However, we open-source the entire analysis package for this Letter, so another group with resources is free to modify, cross-check, and improve upon the posterior computation.

3. Inferences

In this section we report our inferences. We first summarize the measures used to assess model performance. We then discuss an exploratory analysis that examined sensitivity to resolution settings and selected model assumptions, in particular the effects of different assumed atmospheric composition and hot region configuration. We then report high-resolution posterior inferences for the superior model.

3.1. Performance Measures

For pulse-profile modeling with X-PSI a set of performance measures should be considered for each model in the model space, largely following the protocol of Riley et al. (2019). The first measure, given a set of posterior samples, is graphical and the most practical: basic posterior-checking by inspecting for inconsistency between the statistically independent NICER count number variates and the separable sampling distribution from which those variates are assumed to be drawn a posteriori. We estimate the expectation with respect to the posterior of the expected count numbers and form Poisson sampling distributions from these expected count numbers. If there is clear structural difference—namely residual correlations in the joint space of energy and phase—then the model cannot generate data with the structure of the real count numbers. Supposing there are no discernible correlations, then because the sampling distribution for each variate is Poissonian with a sufficiently large expectation for the distribution to be well-approximated as Gaussian, we can inspect the distribution of the standardized residuals to identify any clear deviation from a normal distribution—e.g., too much or too little weight in the tails—

⁶⁰ For additional details about these variants, refer to Riley et al. (2019) and references therein. We also do *not* use the importance sampling algorithm variant (Feroz et al. 2013).

Table 1
Summary Table for ST-U NSX Fully Ionized Hydrogen Hot Regions Plugged into the NICER × XMM Likelihood Function, Conditional on the NANOGrav × CHIME/Pulsar Prior PDF

Parameter	Description	Prior PDF (Density and Support)	$\widehat{\text{CI}}_{68\%}$	\widehat{D}_{KL}	$\widehat{\text{ML}}$
P (ms)	coordinate spin period	$P = 2.8857^{\text{a}}$ fixed
M (M_{\odot})	gravitational mass joint prior PDF $N(\mu^*, \Sigma^*)$	$M, \cos(i) \sim N(\mu^*, \Sigma^*)$ $\mu^* = [2.082, 0.0427]^{\text{T}}$ $\Sigma^* = \begin{bmatrix} 0.0703^2 & 0.0131^2 \\ 0.0131^2 & 0.00304^2 \end{bmatrix}$	$2.072^{+0.067}_{-0.066}$	0.01	2.070
R_{eq} (km)	coordinate equatorial radius compactness condition ^c effective gravity condition ^d	$R_{\text{eq}} \sim U(3r_{\text{g}}(1), 16)^{\text{b}}$ $R_{\text{polar}}/r_{\text{g}}(M) > 3$ $13.7 \leq \log_{10} g_{\text{eff}}(\theta) \leq 15.0, \forall \theta$	$12.39^{+1.30}_{-0.98}$	0.58	11.02
Θ_p (rad)	p region center colatitude	$\cos(\Theta_p) \sim U(-1, 1)$	$1.35^{+0.46}_{-0.39}$	0.25	1.622
Θ_s (rad)	s region center colatitude	$\cos(\Theta_s) \sim U(-1, 1)$	$1.89^{+0.40}_{-0.46}$	0.24	2.303
ϕ_p (cycles)	p region initial phase ^e	$\phi_p \sim U(-0.5, 0.5)$, wrapped ^f	bimodal	3.52	0.185
ϕ_s (cycles)	s region initial phase ^e	$\phi_s \sim U(-0.5, 0.5)$, wrapped	bimodal	3.51	0.243
ζ_p (rad)	p region angular radius	$\zeta_p \sim U(0, \pi/2)$	$0.147^{+0.070}_{-0.041}$	2.18	0.093
ζ_s (rad)	s region angular radius	$\zeta_s \sim U(0, \pi/2)$	$0.146^{+0.071}_{-0.042}$	2.15	0.127
	no region-exchange degeneracy non-overlapping hot regions	$\Theta_s \geq \Theta_p$ function of $(\Theta_p, \Theta_s, \phi_p, \phi_s, \zeta_p, \zeta_s)$			
$\log_{10}(T_p)$ [K]	p region NSX effective temperature	$\log_{10}(T_p) \sim U(5.1, 6.8)$, NSX limits	$5.988^{+0.048}_{-0.059}$	2.95	6.080
$\log_{10}(T_s)$ [K]	s region NSX effective temperature	$\log_{10}(T_s) \sim U(5.1, 6.8)$, NSX limits	$5.992^{+0.047}_{-0.058}$	2.98	6.058
$\cos(i)$	cosine Earth inclination to spin axis	$M, \cos(i) \sim N(\mu^*, \Sigma^*)$	$0.0424^{+0.0029}_{-0.0029}$	0.01	0.044
D (kpc)	Earth distance	$D \sim \text{skewnorm}(1.7, 1.0, 0.23)^{\text{h}}$	$1.21^{+0.15}_{-0.15}$	0.10	0.995
N_{H} [10^{20} cm^{-2}]	interstellar neutral H column density	$N_{\text{H}} \sim U(0, 10)$	$1.587^{+1.953}_{-1.092}$	0.91	0.216
α_{NICER}	NICER effective-area scaling	$\alpha_{\text{NICER}}, \alpha_{\text{XMM}} \sim N(\mu, \Sigma)$	$1.026^{+0.136}_{-0.137}$	0.03	1.111
α_{XMM}	XMM effective-area scaling joint prior PDF $N(\mu, \Sigma)$	$\mu = [1.0, 1.0]^{\text{T}}$ $\Sigma = \begin{bmatrix} 0.150^2 & 0.106^2 \\ 0.106^2 & 0.150^2 \end{bmatrix}$	$0.93^{+0.14}_{-0.13}$	0.17	0.638
	Sampling process information number of free parameters: ^l 15 number of processes: ^j 1 number of live points: 4×10^4 hypervolume expansion factor: 0.1^{-1} termination condition: 10^{-1} evidence: ^k $\ln \mathcal{Z} = -20714.61 \pm 0.02$ number of core ^l hours: 28320 likelihood evaluations: 23710136 nested replacements: 1250386 effective sample size: ^m 420281				

Notes. The description in this caption is largely adopted from Riley et al. (2019) for consistency. We provide: (i) the parameters that constitute the sampling space, with symbols, units, and short descriptions; (ii) any notable derived or fixed parameters; (iii) the joint prior distribution, including hard truncation bounds and constraint equations that define the hyperboundary of the support; (iv) one-dimensional (marginal) 68.3% credible interval estimates symmetric in posterior mass about the median ($\widehat{\text{CI}}_{68\%}$); (v) Kullback–Leibler (KL)-divergence estimates in *bits* (\widehat{D}_{KL}) representing prior-to-posterior information gain (see the appendix of Riley et al. 2019 for a high-level description of the divergence); (vi) the parameter vector ($\widehat{\text{ML}}$) estimated to maximize the background-marginalized likelihood function, corresponding to a nested sample. Note that, strictly, the target of nested sampling is not to generate a maximum likelihood estimator—it is a by-product of the sampling process for evidence estimation. Moreover, there is degeneracy in the likelihood function and high-likelihood solutions with remarkably different parameter values—such as a radius near or above the posterior median—may be retrieved from the public sample information. Constraint equations in terms of two or more parameters result in *marginal* distributions that are not equivalent to those inverse-sampled.

^a Cromartie et al. (2020); Wolff et al. (2021).

^b The function $r_{\text{g}}(M) := GM/c^2$ denotes the gravitational radius with dimensions of length.

^c The coordinate polar radius of the source 2-surface, $R_{\text{polar}}(M, R_{\text{eq}}, \Omega)$, is a quasi-universal function adopted from AlGendy & Morsink (2014), where $\Omega := 2\pi/P$ is the coordinate angular rotation frequency.

^d The range of effective gravity from the equator (minimum gravity) to the pole (maximum gravity) must lie within NSX limits. A quasi-universal function is adopted from AlGendy & Morsink (2014) for effective gravity $g_{\text{eff}}(\theta; M, R_{\text{eq}}, \Omega)$ in units of cm s^{-2} in the table, where $\Omega := 2\pi/P$ is the coordinate angular rotation frequency.

^e With respect to the meridian on which Earth lies.

^f The periodic boundary is admitted and handled by MULTINEST. However, this is an unnecessary measure because we straightforwardly define the mapping from the native sampling space to the space of a phase parameter ϕ such that the likelihood function maxima are not in the vicinity of this boundary.

^g With respect to the meridian on which the Earth antipode lies.

^h Specifically, the PDF defined as `scipy.stats.skewnorm.pdf(D, 1.7, loc = 1.002, scale = 0.227)`, truncated to the interval $D \in [0, 1.7]$ kpc.

ⁱ In the sampling space; the number of background count rate variables is equal to the number of channels defined by the NICER and XMM data sets.

^j The mode-separation MULTINEST variant was deactivated, meaning that isolated modes are not evolved independently and nested sampling threads contact multiple modes. In principle this also allows us to combine the processes in a post-processing phase using *nestcheck* (Higson 2018), if more than one process is available for a given posterior; the posteriors derived in the production analysis are high-resolution, so we neglect combining repeat processes.

^k Defined as the prior predictive probability $p(d_{\text{N}}, d_{\text{X}}, \{\mathcal{B}_{\text{X}}\} | \text{ST} - \text{U})$. Note, however, that in order to complete the reported evidence for comparison to models other than those defined in this work, upper-bounds for the NICER background parameters need to be specified.

^l Intel®Xeon® E5-2697A v4.

^m The effective sample size estimator invoked, following DNest4 (Brewer & Foreman-Mackey 2016, <https://github.com/eggplantbren/DNest4>), is the *perplexity* measure

$$\widehat{\text{ESS}} := \exp\left(-\sum_i^I w_i \log w_i\right),$$

where the $\{w_i\}_{i=1, \dots, I}$ are the sample weights (e.g., Martino et al. 2017).

that would be indicative of noise–model inaccuracies. If this check also passes, then the model has sufficient complexity to generate synthetic data with the structure of the NICER PSR J0740+6620 event data and is adequate for simulation purposes—e.g., for statistical forecasts of the constraining power achievable with future X-ray space telescope concepts such as the enhanced X-ray Timing and Polarimetry mission (Watts et al. 2019; Zhang et al. 2019) and the Spectroscopic Time-Resolving Observatory for Broadband Energy X-rays (Ray et al. 2019).

The second measure we inspect is the maximum likelihood estimate reported by a nested sampling process. Note that nested sampling does not target maximum likelihood estimation, but the drawing of samples from the typical set of a target distribution; the maximum likelihood estimate is therefore also subject to the concentration of prior mass in parameter space. More specifically, we are working with the estimated maximum of a background-marginalized likelihood function given by Equation (5) or one of the likelihood factors (i.e., the NICER or XMM likelihood function). We can use these point estimates to compare, in a simple way, models of the same count number variates. We only graduate to comparison of models using maximum likelihood estimates if the graphical posterior checking described above does not reveal failures. The model that reports the highest maximum likelihood estimate among those that model the same count-number variates has a sampling distribution⁶¹ that captures the most structure in the set of count-number variates. It is plausible, therefore, that a data-generating process defined by that model is the closest approximation of physical reality attained by the models considered. For models that can a posteriori generate data with the structure of the real count-number variates, the maximum likelihood estimate can be used to resolve small differences that human inspection fails to uncover.

The third measure that we examine when comparing models of the same set of count-number variates is the evidence (the prior predictive probability distribution evaluated using the real count-number variates). While maximum likelihood estimates are mere point estimates, the evidence is the expectation of the likelihood function with respect to the joint prior PDF. In one respect, this is powerful because unwarranted prior predictive complexity is penalized: if too much complexity is added to model \mathcal{M} to construct model \mathcal{M}^+ , then for \mathcal{M}^+ the likelihood function over a large swathe of prior mass is smaller than the expected likelihood (with respect to the prior) of model \mathcal{M} , which can entirely negate any localized increases in the likelihood function. In other words, much of the additional complexity is unhelpful because the data generated do not have a similar structure to the real data. On the other hand, penalizing complexity in this way is arguably misleading if the model has phenomenological components: in this Letter the surface hot region models *are* phenomenological.

The evidence, together with a prior mass function of nodes of the model space, may be a biased model selector in our context. Unfortunately, it is necessary in a formal Bayesian framework to use the evidence to marginalize over nodes of the model space in order to compute a posterior PDF of parameters of interest that are shared between nodes. Marginalizing over nodes that differ solely by the atmosphere composition is not problematic. If we do not formally marginalize in such a

manner over all models, however, then we can only report the posterior distribution (marginalized over atmosphere composition) for each hot region model that satisfies the graphical posterior checking criterion, together with the maximum likelihood estimate. The reader is then free to interpret the model-to-model posterior variation as a systematic error estimate by, for instance, weighting the posteriors equally which would roughly lead to credible regions with near-maximum hypervolume (width in one dimension, area in two dimensions, and so on); formally, this is equivalent to defining an implicit prior mass distribution over model space nodes that happens to nullify evidence differences, leading to a uniform posterior mass function of models.⁶² Alternatively, any other weighting can be interpreted as the reader choosing their own prior mass function of model space nodes.

Lastly, when comparing models of the same set of count-number variates, we also consider the tractability of the model. If two models pass graphical posterior predictive checks, and supposing one model is less complex, that model is almost by definition more straightforward to implement and to compute the posterior for. The adequately performing model that requires fewest resources to reproduce or prove erroneous—thereby increasing the robustness and potentially the computation accuracy—can be reasoned to be the most useful in practice.

3.2. Exploratory Analysis

In this section we report on posterior sensitivity to various features of the analysis pipeline. Although one can attempt to probe sensitivity using importance sampling, we opted for nested sampling for every variant of interest. In our sensitivity analyses, our nested sampling resolution settings are lower than for the production analysis because we were ultimately resource-limited. We varied the number of nested sampling live-points and the bounding hypervolume expansion factor; we explored XMM background prior hyperparameter variation; we switched the atmosphere composition from hydrogen to helium; and we varied likelihood function resolution settings. We have not probed sensitivity to event data set selection (namely, NICER detector channel cuts) nor sensitivity to approximation of the atmosphere ionization state as fully ionized.

The posteriors we report in this section were computed using at least 2×10^3 live points and (except where explicitly noted) condition on the ST-U model and either the NICER likelihood function or the NICER and XMM likelihood function. For a full description of this model, and a schematic diagram, see Riley et al. (2019). Briefly, however, ST-U assumes each hot region is a single-temperature spherical cap. The two regions can have completely independent properties (temperature and size) and are free to take any location on the star’s surface provided that they do not overlap (see also the discussion in Section 2.3.1). Our exploratory analysis indicated that this model provided an adequate description of the PSR J0740+6620 data set using the performance measures outlined in Section 3.1. Finally, note that posteriors reported in this section are conditional on a NICER exposure time that was erroneously high by $\sim 2\%$. We corrected this number for a subset of posteriors reported in this section, and for the

⁶¹ Within the continuous set of such distributions associated with the model.

⁶² And more formally still, this would mean the prior mass function is dependent on the data, which is a fallacy.

production analysis (Section 3.3); our posteriors are, however, insensitive to this level of error in exposure time.

3.2.1. Impact of Radio Timing Prior Information

The informative joint NANOGrav and CHIME/Pulsar prior is critical for deriving a useful constraint on the radius of PSR J0740+6620. For comparison, we compute a posterior conditional on a fully ionized hydrogen atmosphere and a diffuse, separable prior PDF of the mass, the distance, and the cosine of inclination angle. The mass prior is such that the joint prior PDF of mass and radius is flat within the prior support (see Table 1). The distance prior PDF is adopted from Igoshev et al. (2016) and displayed in Figure 2, with support $D \in [0.1, 10.0]$ kpc. The prior PDF of the cosine of the inclination angle is isotropic, meaning flat with support $\cos i \in [0, 1]$.

The radius posterior conditional on the diffuse prior is much broader than when we condition on the joint NANOGrav and CHIME/Pulsar prior, and there is no independent indication from the pulse-profile modeling for a high mass (see Figure 5). Once the models are conditioned on the joint NANOGrav and CHIME/Pulsar prior PDF, we gain very little additional information a posteriori from the X-ray likelihood function about the pulsar mass, distance, and inclination angle with respect to the spin axis. We can verify this by examining the marginal posterior distributions in comparison to the respective prior distributions, which is summarized for each parameter by the KL-divergence estimate (Kullback & Leibler 1951).

3.2.2. Impact of X-Ray Telescopes

The NICER likelihood function is sensitive to the basic configuration of the surface hot regions and their temperatures, despite the relatively small number of pulsed counts (those above the phase-invariant background). The XMM likelihood function is less informative both due to the lack of phase information, and because the events are sparse for all three EPIC cameras and moderately consistent with the expected background signal derived from blank-sky exposures. However, the XMM likelihood function acts to reduce the NICER posterior mode volume substantially, affecting the inferred radius and geometry. The reason for this is because the XMM likelihood is sensitive to the combined phase-averaged signal from the hot regions being too bright. Therefore, given the NICER likelihood function, we constrain the contribution to the unpulsed portion of the pulse profile that must be generated by the hot regions rather than the backgrounds. Posterior figures demonstrating this are reserved for Section 3.3.

3.2.3. Effect of Atmospheric Composition

The atmospheric composition of PSR J0740+6620 is not known a priori. We therefore compared $ST-U$ posteriors for hydrogen and helium atmospheres assuming full ionization—see the discussion in Section 2.3.2—in the second online figure of the set associated with Figure 5. The marginal radius posteriors were indistinguishable, although there were some small changes in the properties of the hot regions. However, given the apparent lack of sensitivity to atmospheric composition, inferences reported hereafter are conditioned on a fully ionized hydrogen atmosphere—we do not need to marginalize over the binary atmosphere parameter. Both hot regions are inferred to have effective temperatures $T \approx 10^6$ K, at which partial ionization effects should be small.

3.2.4. Hot Region Complexity

The Riley et al. (2019) analysis of PSR J0030+0451 reported a number of hot region models that provided an adequate description of the data according to their performance measures (largely adopted here). These included $ST-U$ and variants in which one of the hot regions was permitted increasingly complex forms, including rings and crescents. The inferred radius changed as model complexity increased, but evidence calculations showed a substantial improvement in model performance. As a result of this, we deemed the $ST+PST$ model—in which one hot region is a single temperature spherical cap and the other is, a posteriori, a crescent—to be superior for PSR J0030+0451.

For PSR J0740+6620 the $ST-U$ model also provides an adequate description of the data. In order to assess whether additional complexity is useful we also condition on the $ST+PST$ model. The posterior configuration and properties of the hot regions conditional on this more complex model (which includes the possibility of hot regions that are simply spherical caps) did not differ in an important way from the configuration inferred from $ST-U$: the hot region for which more complexity was permissible exhibited degeneracy a posteriori—we were not sensitive to the existence of additional emission structure beyond that of a simple spherical cap, and the evidence estimates are consistent. There was therefore no extended crescent as inferred for PSR J0030+0451; the likelihood function degeneracy included some subset of possible crescent structures—those on smaller angular scales (see Riley et al. 2019 for a discussion about hot region structure degeneracy)—which may be of interest to pulsar modelers. The inferred radius changed very little (see the third online figure of the set associated with Figure 5), and there was no increase in model performance. For this reason we hereafter report inferences exclusively for the $ST-U$ model.

3.2.5. XMM Background Prior Sensitivity

As described in Section 2.5.2, the XMM background is free-form, but each variable (one per channel) has a prior with compact support. For each variable, a flat prior PDF is defined whose width is controlled by a hyperparameter n . For the headline inferences reported in this Letter we used $n=4$, having tested sensitivity to varying n in the range $n \in [0.01, 8]$. In the limit that n tends to zero, the background information would be treated as a point estimate of the XMM background. The posterior distribution of the radius is insensitive to n being varied through the range $n \in [0.01, 4]$; see the fourth online figure of the set associated with Figure 5. It broadens slightly for $n=8$ because fainter combined signals from the hot regions have greater background-marginalized likelihoods, yielding additional posterior weight for higher-radius configurations that reduce the unpulsed component while conserving the pulsed component to satisfy the NICER event data. However, the value $n=8$ is arguably too conservative even when considering potential systematic error.

3.2.6. Likelihood Function Resolution Sensitivity

The X-PSI likelihood function has a number of resolution settings, most notably settings that control the discretization of the computational domain for computation of signals (pulse profiles) incident on telescopes. The photon-specific flux signal we require is an integral over a distant observer’s sky of the

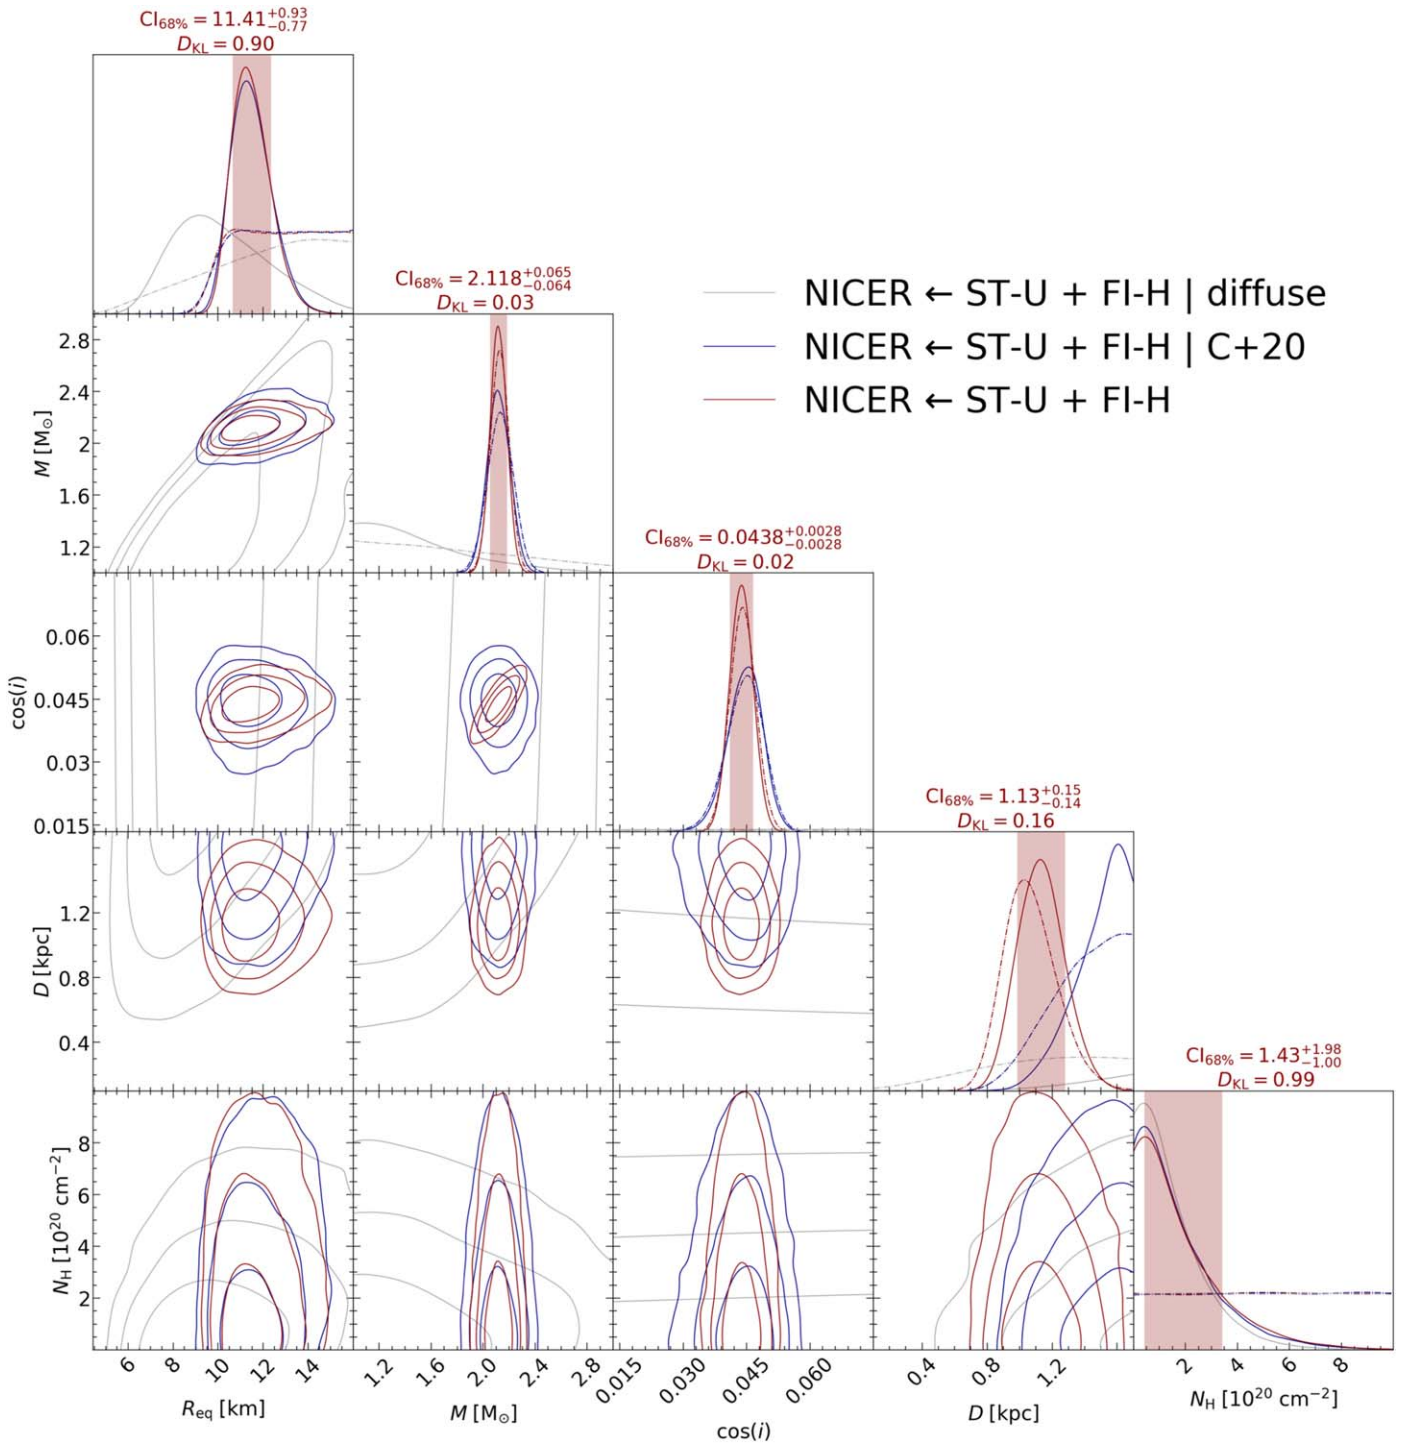


Figure 5. One- and two-dimensional marginal PDFs conditional on the ST-U model, the NICER likelihood function alone, and one of three prior PDFs to probe the impact of radio timing information. From leftmost to rightmost in each panel, the parameters are the equatorial radius, the gravitational mass, the cosine of viewing angle subtended to pulsar spin axis, the distance, and the column density. We display the marginal prior PDFs for each parameter as the dashed-dotted functions; the informative priors encode the information from NANOGraV \times CHIME/Pulsar and Cromartie et al. (2020, denoted by the conditional argument C+20), and the diffuse prior is described in Section 3.2.1. We report estimators for the NICER posterior conditional on the joint NANOGraV and CHIME/Pulsar prior. We report the Kullback–Leibler divergence, D_{KL} , from prior to posterior in *bits* for each parameter. The shaded credible intervals $\text{CI}_{68\%}$ for each parameter are symmetric in marginal posterior mass about the median, containing 68.3% of the mass. The credible regions in the off-diagonal panels, on the other hand, are uniquely the highest-density—and thus the smallest possible—credible regions, containing 68.3%, 95.4%, and 99.7% of the posterior mass. In the appendix of Riley et al. (2019) we provide additional information regarding posterior kernel density estimation (KDE), error analysis, and the estimators displayed here; note that here we use an automated Gaussian KDE bandwidth optimized by GetDist (Lewis 2019). The complete figure set for the exploratory analysis (7 images) is available in the online journal.

(The complete figure set (7 images) is available.)

photon specific intensity from the hot regions, yielding a two-dimensional function of time (rotational phase) and photon energy. The level of discretization with respect to four variables in the domain of the incident photon specific intensity field generally controls the computational expense of likelihood evaluation. These variables are the number of rotational phases and energies the photon-specific flux signal is computed at, the number of hot region surface elements, and the number of rays calculated. See the X-PSI documentation⁶³ for additional information.

We tested posterior sensitivity to increasing the discretization degrees for these variables by recomputing a posterior PDF with a new nested sampling process. We found that doubling these discretization degrees does not yield a change in the posterior PDF that is clearly distinguishable from Monte Carlo sampling noise;⁶⁴ see the fifth online figure of the set associated with Figure 5.

3.2.7. Nested Sampling Resolution Sensitivity

For a fixed bounding hypervolume expansion factor of 0.1^{-1} , the posterior PDFs were insensitive to doubling live-point number from 10^3 up to 4×10^3 . Following sampler comparisons within the NICER collaboration, we then increased resolution to 4×10^4 live points, leading to broadening of the radius posterior; see the sixth and seventh online figures in the set associated with Figure 5. Increasing the sampling resolution by using 8×10^4 live points led to some further broadening, but doubled an already large computational cost. Given the computational resources required for posterior computation with such a large number of live points, we were not able to rigorously prove convergence with live-point number. We decided to adopt 4×10^4 live points for the production analysis; all information necessary to reproduce and improve upon our posterior computation is available in open-source repositories.

Such posterior mode broadening with increased nested sampling resolution is naturally expected because nested sampling algorithms approximate hypersurfaces in parameter space of constant likelihood; these approximations improve with sampling resolution but their sufficiency is difficult to prove for non-trivial likelihood functions encountered in real problems and when subject to resource limitations. It is desirable to transform away nonlinear modal degeneracies so that an approximation conforms more efficiently to structure in the sampling space; however, this can in practice be an intractable task for a given problem. Moreover, when sampling from a target distribution with two or more modes of commensurate posterior mass, the live-point resolution is roughly split between the modes, reducing the resolution of a given mode due to the approximations alluded to above. For PSR J0740+6620, posterior bimodality arises due to the near-equatorial inclination of the source, leading to two competitive geometric configurations of the hot regions (see Section 3.3, where we discuss this further).

3.3. Production Analysis

Our exploratory analysis indicates that model ST-U provides an adequate description of the data and that the posteriors are

⁶³ <https://thomasedwardriley.github.io/xpsi/>

⁶⁴ The nested sampling seed was set based on the system clock for each sampling process and therefore not held constant as would be ideal.




Figure 6. NICER count data $\{d_{ik}\}$, posterior-expected count numbers $\{c_{ik}\}$, and (Poisson) residuals for ST-U. Note that we split the count numbers in the upper two panels over two rotational cycles, such that the information on phase interval $\phi \in [0, 1]$ is identical to that on $\phi \in (0, 2]$; our data-sampling distribution, however, is defined as the (conditional) joint probability of all event data grouped into phase intervals on $\phi \in [0, 1]$. We display the standardized (Poisson) residuals in the bottom panel: the residuals for the rotational cycle $\phi \in [0, 1]$ were calculated in terms of all event data on that interval (as for likelihood definition), and simply cloned onto the interval $\phi \in (1, 2]$. In Section 3.1 and in the appendix of Riley et al. (2019) we elaborate on the information displayed here.

largely insensitive to either atmospheric composition or increased hot region complexity. In this section, we present high-resolution posteriors—using 4×10^4 live points—conditional on the ST-U model, and a fully ionized hydrogen atmosphere. For each posterior we use either the NICER likelihood function alone, the NICER and XMM likelihood function, or the XMM likelihood function alone. The posterior PDF conditional on the NICER likelihood function alone is derived by importance-sampling another posterior PDF, thereby updating a deprecated radio timing prior PDF (see Section 2.2.2 and Fonseca et al. 2021a). The weighted and equally weighted samples from the marginal joint posterior distribution of mass and radius may be found in the persistent repository of Riley et al. (2021), together with credible region contour point-sequences and marginal posterior PDFs of the radius (to facilitate plotting).

Figure 6 provides a simple graphical posterior predictive check on the model performance, demonstrating that the ST-U model can generate synthetic event data that is commensurate with the NICER data. No unexpected structure—such as large deviations or correlations—is emergent in the residuals.

Marginal posterior distributions of the spacetime parameters—in particular the radius—are shown in Figure 7. The figure displays posteriors conditional on the NICER data alone, conditional on the XMM data alone, and conditional on both

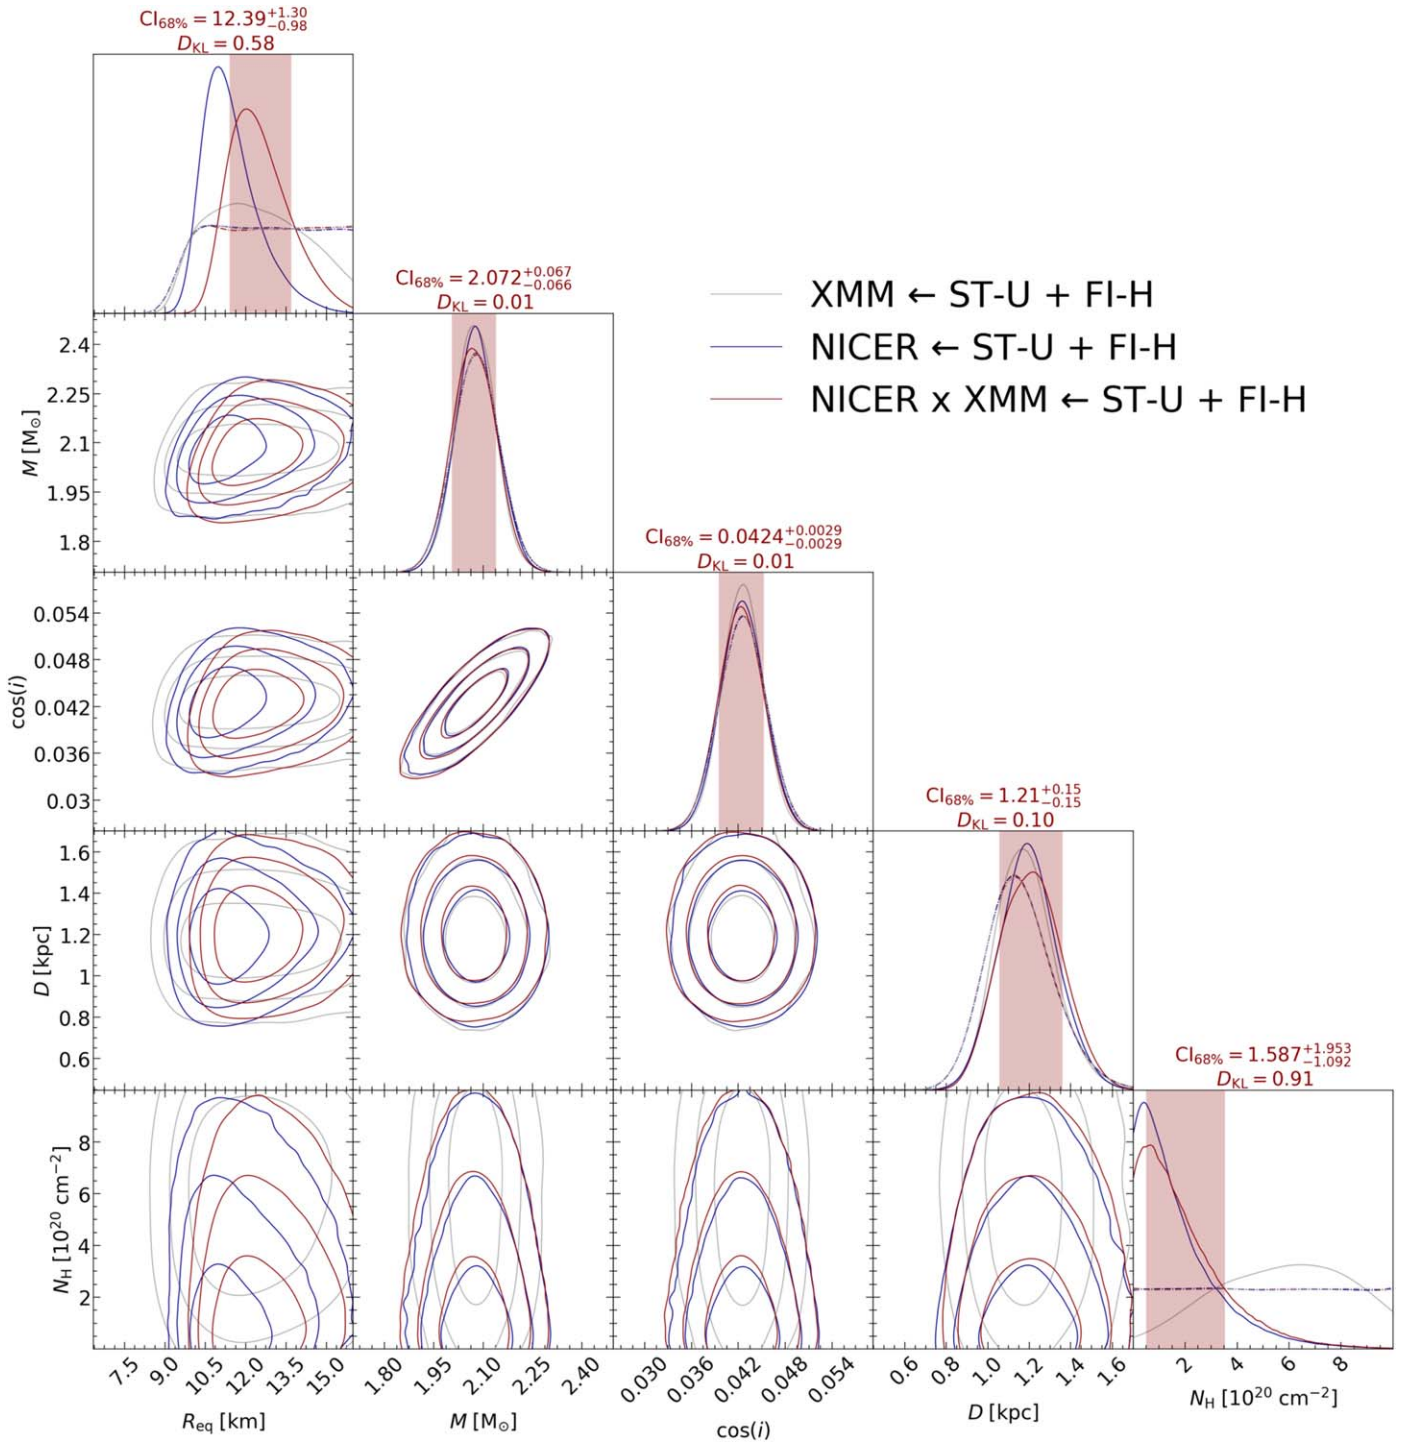


Figure 7. One- and two-dimensional marginal PDFs for fundamental parameters, conditional on the ST-U model and either the XMM likelihood function alone, the NICER likelihood function alone, or the NICER and XMM likelihood function. From leftmost to rightmost in each panel, the parameters are the equatorial radius, the gravitational mass, the cosine of viewing angle subtended to pulsar spin axis, the distance, and the column density. We display the marginal prior PDFs for each parameter as the dashed-dotted functions; the informative priors encode the NANOGrav \times CHIME/Pulsar information. We report estimators for the NICER \times XMM posterior. We use an automated Gaussian kernel density estimation bandwidth optimized by GetDist (Lewis 2019). See the caption of Figure 5 for additional details about the figure elements.

NICER and XMM data. As expected, the phase-resolved NICER likelihood function is far more constraining, in isolation, than the phase-averaged XMM likelihood function. However, the likelihood function product over telescopes excludes regions of the NICER posterior modes where the contribution to the unpulsed component of the pulse profile from the hot regions is too bright. The unpulsed component is

brighter for models in the NICER posterior modes where the star is more compact. Restricting to lower compactness increases the inferred radius for the combined data set by ~ 1 km.

The inferred hot region parameters, again comparing the likelihood functions in isolation to the likelihood function, are shown in Figure 8. The effect of including the XMM data can be

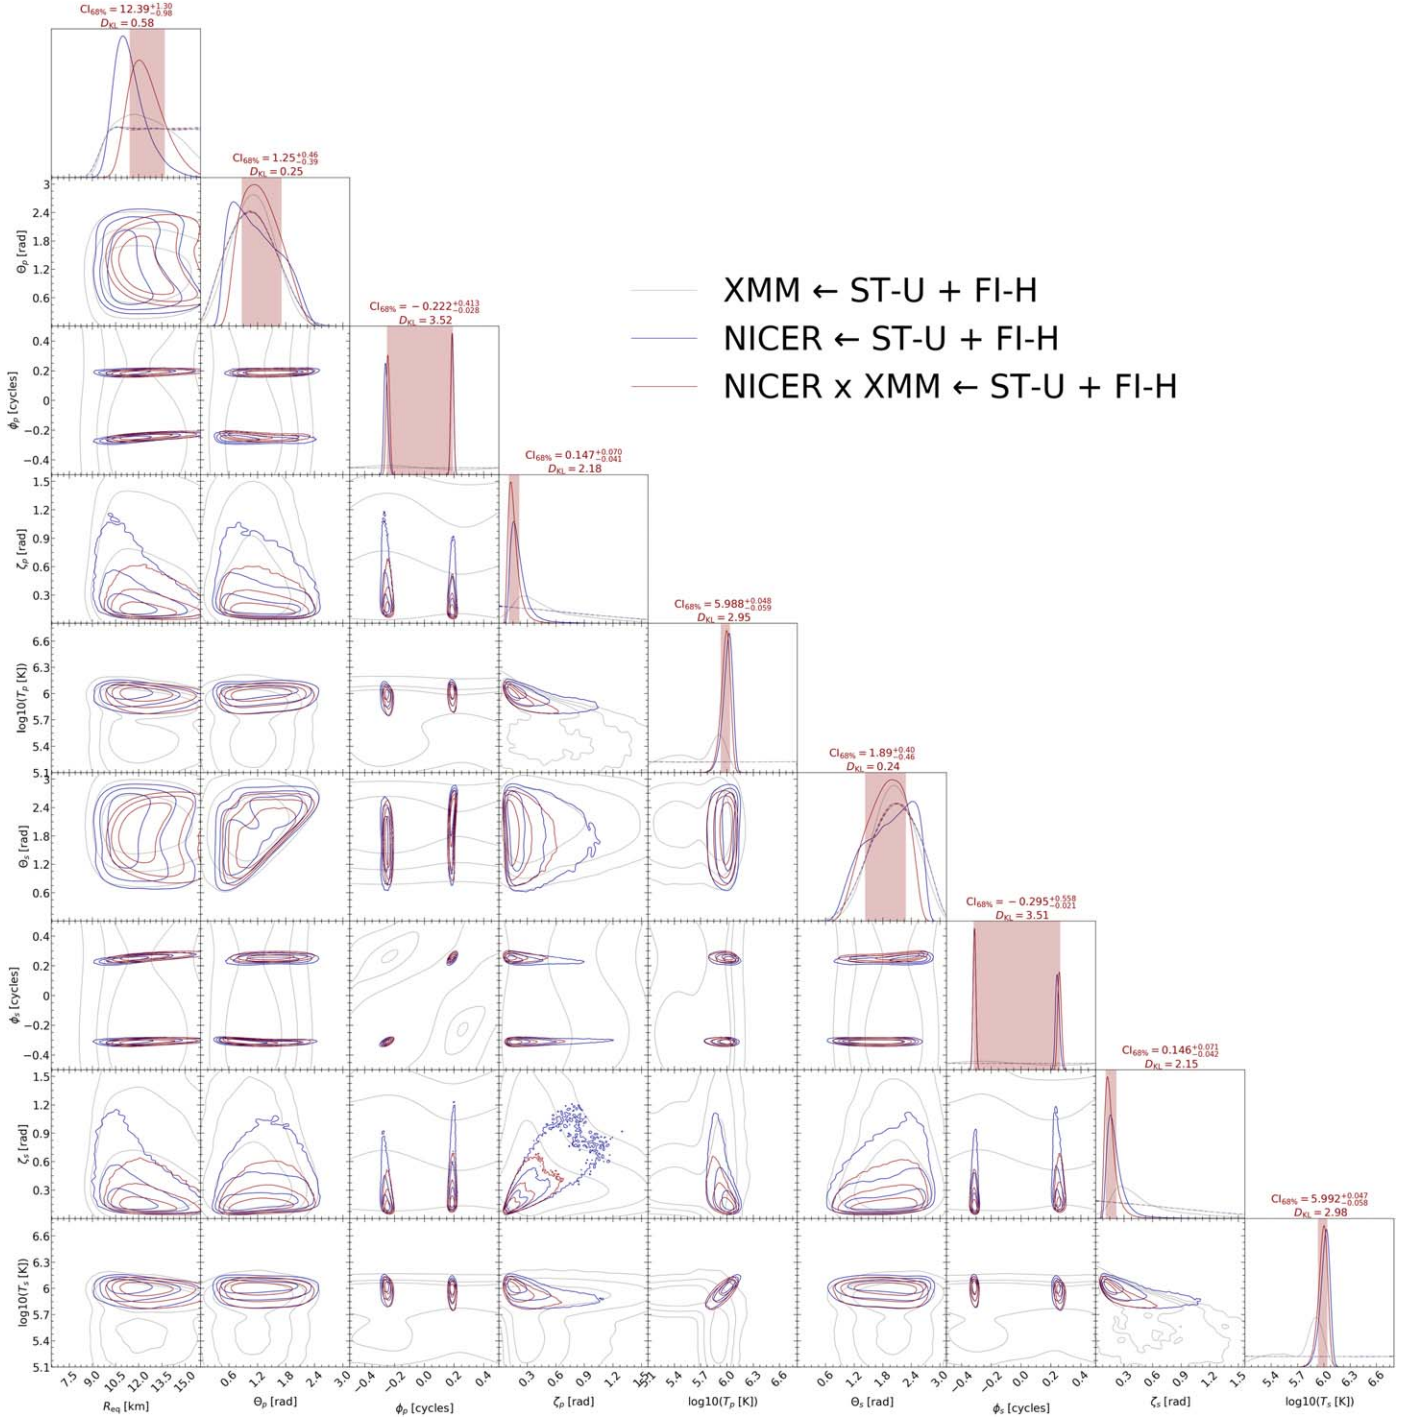


Figure 8. One- and two-dimensional marginal posterior distributions of hot region parameters conditional on the ST-U model. Three types of posterior distribution are rendered: one conditional only on the NICER likelihood function; one conditional only on the XMM likelihood function; and one conditional on the NICER and XMM likelihood function.

seen in the joint posterior PDF of the stellar radius and the angular radii of the two hot regions (ζ_p and ζ_s). The NICER-only posteriors (at the 99.7% level) include models with a smaller stellar radius (<10.5 km) and hot regions with a larger angular radius ($\zeta \sim 1.2$ rad). The large hot regions on very compact stars lead to a bright unpulsed component of the combined signal from those regions. The inclusion of the XMM data means that more of the unpulsed signal is associated with the background instead of the hot regions. As a result, these smaller stars with large hot regions are excluded when the XMM data are included.

Interestingly there are two different posterior modes (due to the near-equatorial inclination),⁶⁵ which can be seen in more detail in the phase-averaged skymaps in Figure 9 and the animated skymap in Figure 10. For neither mode are the two hot regions antipodal. The effect on the inferred signal (pulse-profile and phase-averaged spectrum) of combining the two

⁶⁵ Note that these are different geometric configurations because we expressly exclude hot region exchange degeneracy from the prior support (see Section 2.3.1).




Figure 9. Novel type of figure rendering the phase-averaged expected (photon) specific intensity as a function of sky direction for the source–receiver configuration estimated to maximize the (background-marginalized) $ST-U$ likelihood function given by Equation (5), showing the two different posterior modes and the effect of XMM likelihood function inclusion. Top-left set of four panels: NICER likelihood function—mode one. Bottom-left set of four panels: NICER likelihood function—mode two. Top-right set of four panels: product of NICER and XMM likelihood functions—mode one. Bottom-right set of four panels: product of NICER and XMM likelihood functions—mode two. The expected photon-specific flux spectrum registered by NICER if we phase-average, and (when included) the XMM cameras, is implicitly formed from a fine set of these images. These representative images at four photon energies include all relativistic effects in the likelihood function; note that we extend slightly beyond the XMM waveband used for likelihood evaluation in order to render the (relativistic) rotational effects more vividly. Each panel is normalized to the maximum phase-averaged specific intensity over sky direction at that energy. The background sky has the same intensity—zero—as neighborhoods of the image that a hot region never traverses because the surface exterior of the hot regions is not explicitly radiating in the models (see Section 2.3.3). For animated (photon-) specific intensity sky maps corresponding to these four variants (two posterior mode variants for each of two likelihood function variants), together with pulse-profile traces and spectral evolution, refer to the online journal.

data sets is shown in Figures 11 and 12, respectively. The inclusion of the XMM data reduces the contribution of the hot region emission to the unpulsed component of the pulse profile,

leading to a lower count rate but an increased pulsation amplitude (see the lower panels of Figure 11) in the combined signal from the hot regions.




Figure 10. Summary of the animated figure available in the online journal. In the animated figure, the top three panels show the (photon-) specific intensity as a function of sky direction at three different photon energies as the star rotates. The bottom-left panel displays the (photon-) specific flux pulse profiles traced out by the skymaps, each normalized to its respective maximum. The bottom-right panel displays the (photon-) specific flux spectrum, where the energy bounds each correspond to a skymap energy, as does the vertical line; the trace of the vertical line intersecting the spectrum is one of the pulse profiles. The star rotates 16 times during the 48 s animation. In this summary figure we aim to display the gravitationally lensed geometric configuration of the surface hot regions from our Earthly viewing perspective, over the course of one rotational cycle. We display the (photon-) specific intensity as a function of sky direction at the lowest energy, as the star rotates through the panels from left to right and from top to bottom.

(An animation of this figure is available.)

4. Discussion

4.1. Radius Measurement and Implications for the EOS

The inferred equatorial radius of the massive pulsar PSR J0740+6620 is $R_{\text{eq}} = 12.39^{+1.30}_{-0.98}$ km, where the credible interval bounds are approximately the 16% and 84% quantiles

in marginal posterior mass, given relative to the median. The inferred mass, $2.072^{+0.067}_{-0.066} M_{\odot}$ is dominated by the mass prior from the radio timing, $2.08 \pm 0.07 M_{\odot}$. The 90% credible interval for the radius is $12.39^{+2.22}_{-1.50}$ km and the 95% credible interval is $12.39^{+2.63}_{-1.68}$ km. It is worth stressing that when we carried out pulse-profile modeling without using the mass prior




Figure 11. Posterior-expected pulse profiles conditional on the ST-U hot region model and either the NICER likelihood function (left panels) or the NICER and XMM likelihood function (right panels). We show the signal incident on the telescopes (top and top-center panels) and as registered by NICER (bottom-center and bottom panels). The signal in the top panels has been integrated over the linearly spaced instrument energy intervals, and is effectively proportional to the photon-specific flux. The black curves are the posterior-expected signals generated by the hot regions in combination. We also represent the conditional posterior distribution of the incident photon flux (top-center panels) and the NICER count rate (bottom panels) at each phase as a set of one-dimensional highest-density credible intervals, and connect these intervals over phase via the contours; these distributions are denoted by π (photons $\text{cm}^{-2} \text{s}^{-1}; \phi$) and π (counts $\text{s}^{-1}; \phi$). Note that the fractional width of the credible interval at each phase is usually higher for π (photons $\text{cm}^{-2} \text{s}^{-1}; \phi$) than for π (counts $\text{s}^{-1}; \phi$) because of the variation permitted for the instrument model; in combination, the signal registered by the instrument is more tightly constrained. To generate the conditional posterior bands we apply the X-PSI package, which in turn wraps the `fgivenx` (Handley 2018) package.

from radio timing, we would not have inferred independently from the NICER and XMM data alone that PSR J0740+6620 is a high-mass source (nor did we obtain any informative constraint on the radius; see Section 3.2.1).

Rotation increases the equatorial radius of a neutron star. The increase in radius for fixed mass is small, as shown in Figure 13, where mass–radius curves are plotted for four representative EOSs that span a wide range of allowed stiffness (Hebeler et al. 2013). Mass–radius curves for non-rotating stars and stars rotating at 346 Hz are shown for each EOS. For a $2.0 M_{\odot}$ star, the equatorial radius increases by slightly more than 0.05 km for the softest EOS, while the increase is as large as 0.2 km for the stiffest EOS. These increases in radius due to spin are smaller than our uncertainty in measuring the neutron star’s radius at present. The change in radius due to spin is already incorporated into our pulse-profile models, since we

assume the rotating star is deformed into an oblate shape. While the shape function that we use is an approximation, Silva et al. (2021) have shown that it is sufficiently accurate for all of the rotation-powered pulsars with spin frequencies less than 400 Hz that NICER observes.

The radius inferred for PSR J0740+6620 is very similar to that inferred from pulse-profile modeling of NICER data for PSR J0030+0451, although for the latter the inferred mass was lower: Riley et al. (2019) found $R_{\text{eq}} = 12.71^{+1.14}_{-1.19}$ km and $M = 1.34^{+0.15}_{-0.16} M_{\odot}$; the independent analysis of Miller et al. (2019) found $R_{\text{eq}} = 13.02^{+1.24}_{-1.06}$ km and $M = 1.44^{+0.15}_{-0.14} M_{\odot}$. Note that the width of the credible interval on the radius for PSR J0740+6620 is not smaller, despite the constraining mass prior; this is due to the lower number of source counts, which limits the precision of the radius measurement. The radius that




Figure 12. Posterior-expected spectra conditional on the ST-U hot region model and either the NICER likelihood function (left panels) or the NICER and XMM likelihood function (right panels). We show the spectrum that would be incident on the telescopes if it were unattenuated by the interstellar medium (top panels) and the signal as registered by NICER (center and bottom panels). The black rate curves are the posterior-expected spectra generated by the hot regions in combination. We represent the conditional posterior distribution $\pi(\text{photons keV}^{-1} \text{ cm}^{-2} \text{ s}^{-1}; E)$ of the unattenuated incident photon-specific flux at each energy as a set of one-dimensional highest-density credible intervals, and connect these intervals over phase via the contours (top panels); the energies displayed are those spanning the waveband of the NICER channel subset [30, 150]. In the center panels we display the background-marginalized posterior-expectation of the source count-rate signals, *plus* the background count-rate terms that maximize the *conditional* likelihood functions; the center-right signal is equivalent to that displayed in the center panel of Figure 6. In the bottom panels we display the posterior-expected count-rate spectra generated by the hot regions in combination and individually, together with the conditional posterior NICER count rate distribution $\pi(\text{counts s}^{-1}; \text{channel})$ for each channel. Moreover, the topmost green step functions are the phase-average of the center panels—each is effectively, but not exactly, the observed count-number spectrum divided by the total NICER exposure time.

we report for PSR J0740+6620 is also consistent with the values inferred from the most recent phase-averaged spectral modeling of quiescent and bursting neutron stars (Nättilä et al. 2017; Steiner et al. 2018; Baillot d’Etiavaux et al. 2019; González-Caniulef et al. 2019, noting that for some of these analyses a neutron star mass is assumed rather than inferred or known in advance), and with indirect constraints from the inner radii of accretion disks (Ludlam et al. 2017).

The detection of gravitational waves from binary neutron star mergers provides an alternative method of constraining the dense matter EOS. A measurement of tidal deformability and neutron star masses from the late inspiral phase can be used to infer EOS parameters and hence the associated mass–radius relation. The posteriors on any mass–radius relation derived in this way depend on the EOS model and the priors on the model parameters, and this should be kept in mind when comparing them to the radii inferred directly (without reference to EOS models) from pulse-profile modeling (Greif et al. 2019). Nevertheless the radii derived from the tidal deformability of the binary neutron star merger GW170817 are (for a range of EOS models) lower than the value we derived for PSR J0740+6620 (see for example Abbott et al. 2018, 2019; De et al. 2018; Most et al. 2018; Essick et al. 2020; Landry et al. 2020;

Li et al. 2020). However, the credible intervals on the gravitational-wave-derived radii are of similar extent, and thus the results are certainly consistent with those derived by NICER.

As is clear from Figure 13, the radius of PSR J0740+6620 is in the center of the range considered plausible for neutron stars, $\sim 2M_{\odot}$, and appears compatible with a wide range of EOS models (see, e.g., Hebeler et al. 2013; Greif et al. 2019). However, full Bayesian inference of EOS models is required to fully quantify the constraints arising. For this we refer the reader to the companion paper by Raaijmakers et al. (2021), which carries out EOS inference using results from NICER both individually and in combination with constraints from gravitational-wave observations and their electromagnetic counterparts. Using two different high-density EOS parameterizations, and models that connect to microscopic calculations of neutron matter from chiral effective field theory interactions at nuclear densities, Raaijmakers et al. (2021) show that the new NICER results provide tight constraints, for example on the pressure of neutron star matter at around twice saturation density.

The measurement of radius for a high-mass neutron star is also interesting for the properties of potential quark cores (see




Figure 13. Mass vs. equatorial radius for several example EOS models from Hebeler et al. (2013), showing the difference between non-rotating stellar models and stars rotating at 346 Hz. For each EOS shown the right hand (heavier) curve is for a spin of 346 Hz, while the left-hand (lighter) curve is for zero rotation. The 68% and 95% credible regions for mass and radius inferred from our analysis of PSR J0740+6620 are shown by the shaded cyan contours. The blue crosshair shows the inferred median values.

for example Annala et al. 2020). Han & Prakash (2020) consider the implications of such a measurement for our understanding of quark matter phases in neutron stars for different model types: self-bound strange quark star models, hybrid star models with different types of phase transition, and third-family models where two branches with different radii are possible for the same mass (Schertler et al. 2000). Our upper and lower bounds on the radius posterior at high mass disfavor some regions of quark matter parameter space: both the stiffness of the strange quark phase and the transition properties. The inferences are moreover quite restrictive for self-bound strange quark stars and third-family stars, both of which typically have low radii at high mass.

We can also look at the implications of the change in radius as one moves from $M \sim 2.0 M_{\odot}$ to $M \sim 1.4 M_{\odot}$, an important distinguishing characteristic of different EOS models (Greif et al. 2019; Han & Prakash 2020; Xie & Li 2020; Drischler et al. 2021). Generally, hadronic EOSs having symmetry energy parameters in the ranges predicted by nuclear mass fits and neutron matter studies and with $M_{\max} \lesssim 2.2 M_{\odot}$ would result in $\Delta R = R_{2.0} - R_{1.4} \lesssim -1$ km. The above studies show that matter with a phase transition around $2n_{\text{sat}}$ to a relatively soft phase with sound speed squared $c_s^2 \sim 1/3$ (such as to non-interacting quark matter) would also result in $\Delta R \lesssim -1$ km. Such models also have $M_{\max} \lesssim 2.2 M_{\odot}$. In contrast, larger values of $\Delta R \lesssim 0.5$ km suggest either stiffer high-density matter without a phase transition having $M_{\max} \sim 2.3\text{--}2.5 M_{\odot}$, or a transition to a relatively stiff phase at a transition density $\geq 2.6n_{\text{sat}}$ (Drischler et al. 2021). Even larger values of $\Delta R > 0.5$ km would imply a transition at a lower density $\leq 2n_{\text{sat}}$ to similarly stiff matter. The companion paper by Raaijmakers et al. (2021), which utilizes parameterized EOS models constrained by theories of neutron matter, together with observations of pulsar masses, gravitational waves from mergers, and the X-PSI NICER results for PSR J0030+0451 and PSR J0740+6620, infers that $\Delta R \simeq -0.5^{+1.2}_{-1.5}$ km averaged

over EOS models. This is consistent with the direct observational value $R_{J0740} - R_{J0030} = -0.3^{+1.2}_{-1.5}$ km (this paper; Riley et al. 2019). Although values of $\Delta R < -1$ km and $\Delta R > +1$ km cannot be ruled out, these results suggest more moderate values of ΔR that favor relatively stiff dense matter with a large M_{\max} or an EOS with stiffening at a density $2\text{--}3n_{\text{sat}}$, which could result from a first-order phase transition or a crossover transition like that due to the appearance of quarkyonic matter (McLerran & Reddy 2019). Observational upper limits to M_{\max} , such as the value $M_{\max} \lesssim 2.2\text{--}2.3 M_{\odot}$ suggested by GW170817 (Margalit & Metzger 2017), could help distinguish these possibilities.

The maximum mass of neutron stars, and hence the boundary between the neutron star and black hole populations, is also a function of the EOS. Currently feasible EOS models would permit a maximum neutron star mass in the range $2\text{--}3 M_{\odot}$, but stiffer EOS, with larger radii, are required to achieve higher masses. Analysis of the electromagnetic counterpart of the binary neutron star merger GW170817 has suggested a maximum neutron star mass somewhere in the range $2.0\text{--}2.3 M_{\odot}$ (Margalit & Metzger 2017; Rezzolla et al. 2018; Ruiz et al. 2018; Shibata et al. 2019), but there is a strong dependence on how the kilonova is modeled. Nevertheless this range is consistent with the relatively soft EOS inferred from the tidal deformability for GW170817 (see above). The recent detection of GW190814, a binary compact object merger involving an object with mass $\sim 2.6 M_{\odot}$ (Abbott et al. 2020), is however intriguing. There is considerable debate over whether this object could be a high-mass neutron star rather than a low-mass black hole (Fattoyev et al. 2020; Huang et al. 2020; Sedrakian et al. 2020; Tan et al. 2020; Tsokaros et al. 2020; Zhang & Li 2020b; Dexheimer et al. 2021; Drischler et al. 2021; Godzieba et al. 2021; Tews et al. 2021) and still be consistent with GW170817. The radius that we have inferred for PSR J0740+6620 suggests a lower maximum mass, however (see Raaijmakers et al. 2021).

4.2. Constraining Power of XMM

The likelihood function given NICER and XMM data sets is dominated by the information from the former. However, longer XMM exposure times naturally yield greater constraining power. A deep exposure exists for the rotation-powered millisecond PSR J0030+0451 (Bogdanov & Grindlay 2009) and, as suggested by Riley et al. (2019), the associated spectroscopic (and timing) event data can be jointly modeled with the NICER event data set to address the open question regarding the contribution of the model hot regions to that set of events, which Miller et al. (2019) and Riley et al. (2019) inferred to be (close to) minimal. The term *minimal* is used to mean that the hot regions contributed *only* to the pulsed component of the pulse profile and did not contribute to the unpulsed component. This Letter offers a demonstration of how this can be executed, albeit with a contribution from XMM that is less informative than the contribution from NICER.

The XMM data set for PSR J0740+6620 is phase-averaged and sparse in terms of overall counts, which renders it less constraining than the NICER data set. However, being an imaging telescope, XMM facilitates better quantification of the contribution from the star (attributed to hot regions in our models) compared to the background, whereas the NICER background is more difficult to constrain both a priori and a posteriori. The contribution of the hot regions to the unpulsed

component of the NICER pulse-profile is constrained by jointly modeling the NICER and XMM event data.

For PSR J0740+6620, the contribution from the hot regions is not minimal. Even when one considers only the NICER data set, the hot regions are inferred to generate not only the pulsed component but also part of the unpulsed component of the pulse profile. The effect of including the XMM data in the analysis is to increase the amplitude of the emission from the hot regions by reducing the contribution of the hot regions to the unpulsed component of the pulse profile. Smaller radius stars have larger gravitational fields and cause stronger gravitational lensing. The lensing makes it possible to see the hot regions for a larger fraction of the spin period; the resulting signal has a lower pulsed fraction than the signal from a larger star with less lensing. The samples where the NICER-only unpulsed signal is brighter are those where PSR J0740+6620 is more compact; by weighting away from these, the inclusion of XMM data pushes the posterior toward less compact stars, where the radius is higher (see also the discussion in Miller 2016).

An interesting question is whether the analysis presented in this Letter tells us anything about the effect that a full joint inference analysis might have on the radius inferred for PSR J0030+0451. The emission from the hot regions for PSR J0030+0451 conditional on only NICER event data was minimal, meaning the unpulsed component of the combined signal from the hot regions was small. It follows that the inclusion of the XMM constraints can only increase the contribution from the hot regions. However, the magnitude of the increase in brightness and the effect on the inferred radius is hard to predict because several parameters in the model are degenerate and changes in radius can be offset by, e.g., changes in hot region parameters.

Finally, we explore sensitivity to prior information about the NICER and XMM energy-independent effective area scaling factors α_{XTI} and α_{pn} respectively, and the equatorial radius. The XMM scaling factor α_{XMM} is shared by the three cameras and is denoted by α_{XMM} in Table 1. The blue NICER \times XMM posterior is shown in Figure 7 as the headline posterior; the properties of this posterior are reported in Table 1. The red NICER \times XMM|EAPC is conditional on a compressed effective area prior that aims to approximate the telescope calibration uncertainties discussed in Section 2.4; the acronym EAPC simply means *effective area prior compression*. The prior PDFs are displayed as the dashed-dotted distributions in each on-diagonal panel.




Figure 14. One- and two-dimensional marginal posterior distributions of the NICER and XMM energy-independent effective area scaling factors α_{XTI} and α_{pn} respectively, and the equatorial radius. The XMM scaling factor α_{XMM} is shared by the three cameras and is denoted by α_{XMM} in Table 1. The blue NICER \times XMM posterior is shown in Figure 7 as the headline posterior; the properties of this posterior are reported in Table 1. The red NICER \times XMM|EAPC is conditional on a compressed effective area prior that aims to approximate the telescope calibration uncertainties discussed in Section 2.4; the acronym EAPC simply means *effective area prior compression*. The prior PDFs are displayed as the dashed-dotted distributions in each on-diagonal panel.

instrument cross-calibration is an important aspect of these analyses that warrants careful treatment.

Obtaining estimates of the absolute effective area of an X-ray instrument is a challenging task. Cross-calibration efforts by the International Astrophysical Consortium for High Energy Calibration using observations from multiple concurrent X-ray telescopes have found offsets typically within $\pm 10\%$ but with occasional discrepancies reaching up to $\sim 20\%$ (see, e.g., Ishida et al. 2011; Madsen et al. 2017; Plucinsky et al. 2017). The tails of the joint posterior PDF of the effective scaling parameters (see Figure 14) go beyond what these calibration measurements indicate, so the resulting radius credible intervals should be taken as conservative estimates.

4.2.1. NICER and XMM Backgrounds

The NICER background is difficult to estimate directly, but there are two tools available. Figure 15 shows the NICER background estimated using the “space weather” model (Gendreau 2020) and the “3C50” model (Remillard et al. 2021); see also Bogdanov et al. (2019a). The former models background due to the space weather environment, which varies as NICER moves through different geomagnetic latitudes, and depends on solar activity. The “3C50” model is empirical, taking into account two different types of particle-induced events, the cosmic X-ray background, and a soft X-ray noise component related to operation in sunlight. We also render a set of NICER background estimates for comparison: using joint NICER and XMM posterior samples, we display in

⁶⁶ Such a small shift is not expected to have any remarkable effect on EOS inference, given typical EOS priors (Greif et al. 2019). This is demonstrated in Raaijmakers et al. (2021), where EOS inference is carried out using both our NICER-only inferred radius and our NICER \times XMM inferred radius. Despite an overall change in the median radius posterior inferred from the pulse profile modeling $\sim +1.1$ km once the XMM data set is included, the mass–radius band shifts by a much smaller amount than this, due to the strong influence of the priors on the EOS model.




Figure 15. Comparison of the NICER background conditional on the NICER likelihood function to that conditional on the NICER and XMM likelihood function. The blue step function is the total NICER count spectrum, assumed to be generated by the surface hot regions and the phase-invariant background in our modeling. The solid black step function is the background that maximizes the conditional likelihood function given the parameter vector associated with the nested sample reporting the highest value of the background-marginalized likelihood function. The orange step functions (of which there are 10^3) that form a band are defined similarly, but each is conditional on a sample from the joint NICER and XMM posterior. To ensure tractability, we marginalize over background parameters in order to define a sampling space with $\mathcal{O}(10)$ dimensions; it follows that we cannot estimate marginal posterior PDFs from our posterior samples for each background variable due to information loss. Strictly, the orange band should therefore *not* be interpreted as a collection of posterior PDFs—one per background variable—but as indicative of where posterior mass will be concentrated. We compare the backgrounds to estimates of the NICER background generated using the NICER “space weather” background estimation tool (Gendreau 2020) and the “3C50” model (Remillard et al. 2021). We provide a supplementary figure that shows the NICER background for the 10^3 highest-likelihood nested samples, given the NICER and XMM likelihood function; we also provide a figure that shows the NICER background for a set of 10^3 posterior samples after compression of the joint prior PDF of the telescope effective areas (see Figure 14 and associated text). The complete figure set (3 images) is available in the online journal.

(The complete figure set (3 images) is available.)

Figure 15 the NICER background spectrum that maximizes the conditional likelihood function, yielding a band that is a proxy for background variable posterior mass.

The background spectra displayed in Figure 15 exceed the space weather model at low energies. This excess is not unreasonable given the presence of multiple other point sources in the NICER field of view.⁶⁷ In higher channels, the background spectra appear to agree well with the space weather model. There is a possible small systematic under-prediction compared to the space weather model in channels 80–100; the level of agreement is satisfactory given the current

⁶⁷ The NICER background variables in principle also capture phase-invariant emission from the environment of PSR J0740+6620 that does not originate from the surface hot regions in the X-PSI model. The XMM background prior information is conservative (see Section 2.5.2) and can also in principle capture such phase-invariant emission. The XMM likelihood function is not purely marginalized with respect to an informative blank-sky background prior, which could attribute too much emission to the hot regions in the X-PSI model.

uncertainties on the background modeling. The level of agreement with the 3C50 model, which exceeds the space weather model at low energies and is consistent with it at higher energies, also appears good. Neither background model includes off-axis X-ray sources that might contaminate the target spectrum, and therefore inferred backgrounds that exceed the two estimates are not in principle a problem.

Once the uncertainties on the two NICER background models are understood more fully, we anticipate being able to use them to reduce systematic error in the pulse-profile modeling. The current radius measurement conditional on the XMM data set—which yields an indirect NICER background constraint—will therefore be superseded. Efforts are also ongoing to quantify the level of background due to any off-axis sources in the field of view.

In Figure 16 we display the XMM pn background spectra that maximize the conditional likelihood function given nested samples from the joint NICER and XMM posterior, together with supplementary information. Graphical checking of the spectra against the blank-sky estimate and the event data does not reveal any problems a posteriori.

4.3. Hot Region Configuration

The hot region configuration is assumed to be related to the star’s magnetic field structure. In our previous analysis of PSR J0030+0451, the superior model was $ST+PST$: one of the hot regions was a small spherical cap and the other a long extended arc. A configuration in which the hot regions were antipodal was strongly disfavored. Although the hot regions were separated by approximately 180° in longitude, both were in the same hemisphere of the star.

An antipodal configuration is also disfavored for PSR J0740+6620, once again arguing against a simple dipolar model (although the configuration is closer to antipodal than it is for PSR J0030+0451). The location of the emitting regions is, however, rather different. The expected number of counts contributing to the total expected NICER signal for PSR J0740+6620 is not (close to) minimal a posteriori, despite the diffuse XMM likelihood function (Section 4.2). Only one of the hot regions vanishes from sight during the rotational cycle; the other remains visible at all times. For PSR J0030+0451, the expected number of counts contributing to the total expected NICER signal is (close to) minimal a posteriori for all models that passed graphical posterior predictive checking (Riley et al. 2019); in all cases the hot regions were inferred to dance around the stellar limb, each being entirely non-visible for a substantial fraction of a rotational cycle.

We considered a range of shapes for the hot regions, from circles to rings and arcs. For PSR J0740+6620 the $ST-U$ model (in which both hot regions are circles) provides a reasonable description of the data in terms of, e.g., residuals. A more complex model, $ST+PST$ (the superior model for PSR J0030+0451) did not offer any improvement in model quality measures nor did it lead to changes in the inferred radius or hot region geometry. For PSR J0030+0451 $ST-U$ provided a reasonable description of the data, but we were sensitive a posteriori to additional complexity in the structure of one of the hot regions; consequently, a large shift in the inferred radius was reported. No extended arc structure is inferred for the secondary hot region for PSR J0740+6620, although the hot region could well be a ring instead of a spherical cap.

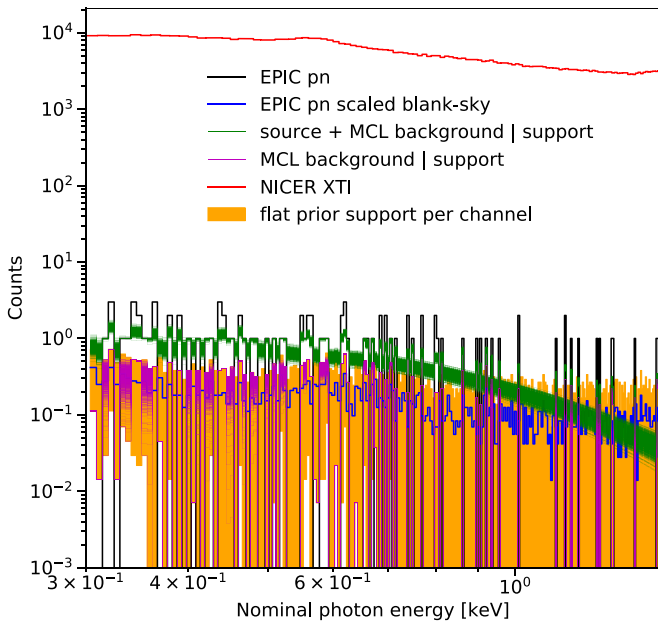


Figure 16. XMM pn background spectra (added to Figure 4) that maximize the conditional likelihood function (denoted by MCL in the legend) given nested samples from the joint NICER and XMM posterior, but subject to the prior support of the background variables. That is, if the maximum of the conditional likelihood function is not within the prior support, the nearest value of the background to the maximum that is within the prior support is used in the spectrum. We also show the total spectra as the sum of the XMM pn source spectra (given the nested samples from the joint NICER and XMM posterior) and the XMM pn background spectra that maximize the conditional likelihood function. The complete figure set (3 images) for the three XMM cameras is available in the online journal.

(The complete figure set (3 images) is available.)

The pulse profile modeling presented in this work constrains the location and shape of the hot regions on the neutron star surface. These regions arise either via heating by magnetospheric currents (Kalapotharakos et al. 2021), or through complex magneto-thermal evolution in the stellar crust (De Grandis et al. 2020). Thus, the information obtained can be used as input for modeling magnetic field structure both in the magnetosphere and inside the star; however, there are currently many unknowns in the picture.

Qualitatively, the pulsars that spin faster have more compact magnetospheres and larger (and more complex if the field has a substantial non-dipolar component) open field line regions. If heating happens at the open field line footprints, then one would expect heated regions of PSR J0740+6620 (with the ratio of light cylinder to neutron star radii, $R_{LC}/R_{NS} = 11.1$) to be larger than those of PSR J0030+0451 ($R_{LC}/R_{NS} = 18.3$), provided that both pulsars have field configurations of similar complexity (i.e., similar relative magnitude of higher-order components), contrary to what is being inferred from the data.

Detailed modeling of pulsar magnetic fields similar to that performed by Kalapotharakos et al. (2021), together with an analysis of crustal thermal evolution, would be interesting from an evolutionary point of view. PSR J0740+6620 has a white dwarf companion while PSR J0030+0451 is solitary and the difference in recent accretion history may play a role in field configuration and residual heating pattern. Their masses are also substantially different, and according to the population study by Antoniadis et al. (2016), such a large difference

cannot be attributed to accretion alone and must stem partly from the difference in progenitor properties.

4.4. Analysis Cross-check

An independent analysis carried out within the NICER collaboration by Miller et al. (2021) reports a PSR J0740+6620 radius of $13.71^{+2.62}_{-1.50}$ km, derived from their combined NICER and XMM analysis. The 68% credible intervals overlap with those that we report in this Letter, but the differences deserve some explanation. Recall that our pulse-profile modeling involves several elements: the NICER phase-resolved data set, the XMM phase-averaged data set, a model for the generation of the count data (including priors on the model parameters), and statistical samplers.

Let us first focus on the analysis of the NICER data. The two teams make different choices on what energy channels to include in the NICER data set: we use channels [30, 150] whereas Miller et al. (2021) use channels [30, 123] (although Miller et al. report that including higher channels does not lead to notable changes in their results).

The two teams also make a number of different prior choices. Miller et al. (2021) assume priors on the mass, distance, and inclination with larger spread to account for potential systematic error on top of the values reported by Fonseca et al. (2021a). Miller et al., unlike us, do not impose a hard upper-limit on the prior support of the radius (see Section 2.2.1): they assume a flat prior on the reciprocal of the compactness $R_{eq}/r_g(M) \sim U(3.2, 8.0)$.⁶⁸ We define the prior support so as to exclude hot-region exchange degeneracy—thus halving the number of posterior modes—whereas Miller et al. do not exclude exchange degeneracy. We also condition on different prior PDFs of the hot region center colatitudes and effective temperatures: the prior PDFs of our hot region center colatitudes are isotropic,⁶⁹ and our prior PDFs of the logarithms of the effective temperatures are uniform. Finally, we use a marginal prior distribution of the energy-independent NICER effective area scaling factor that has a larger spread and broader prior support than Miller et al.; as we discuss below, however, this is not thought to be important. Our prior PDFs are defined in Table 1.

The two teams implement different statistical sampling protocols. We use nested sampling (MultiNest) with high-resolution settings, while Miller et al. (2021) use a hybrid nested sampling (MultiNest) and parallel-tempering ensemble Markov chain Monte Carlo scheme; Miller et al. use far more core hours during the ensemble sampling phase of their computations than during the nested sampling phase. Where both teams use MultiNest, our resolution settings are higher:⁷⁰ we use up to 8×10^4 live-points with a bounding hypervolume expansion factor of 0.1^{-1} , and we eliminate hot-region exchange degeneracy. For the same target distribution, using a higher number of live points and eliminating hot-region exchange degeneracy (and thus the halving the number of modes) both yield lower likelihood—isosurface bounding

⁶⁸ The absence of prior support for high radii is effectively incorporated at a later stage, in the EOS analysis carried out by Miller et al. (2021).

⁶⁹ Meaning uniform in the cosine of the hot region center colatitude.

⁷⁰ A caveat is that the performance of nested sampling with given resolution settings, on the same target distribution, is also dependent on the structure of the likelihood function in the native sampling space—the native sampling space is not, however, unique because different transformations can be defined to inverse-sample a particular joint prior PDF.

Table 2Summary Table for $ST-U$ NSX Fully Ionized Hydrogen Hot Regions Plugged into the NICER Likelihood Function, Conditional on the NANOGrav \times CHIME/Pulsar Prior PDF

Parameter	Description	Prior PDF (density and support)	$\widehat{CI}_{68\%}$	\widehat{D}_{KL}	\widehat{ML}
P (ms)	coordinate spin period	$P = 2.8857$, fixed
M (M_{\odot})	gravitational mass	$M, \cos(i) \sim N(\mu, \Sigma)$	$2.078_{-0.063}^{+0.066}$	0.01	2.125
	joint prior PDF $N(\mu^*, \Sigma^*)$	$\mu^* = [2.082, 0.0427]^T$ $\Sigma^* = \begin{bmatrix} 0.0703^2 & 0.0131^2 \\ 0.0131^2 & 0.00304^2 \end{bmatrix}$			
R_{eq} (km)	coordinate equatorial radius	$R_{eq} \sim U(3r_g(1), 16)$	$11.29_{-0.81}^{+1.20}$	0.72	10.90
	compactness condition	$13.7 \leq \log_{10} g_{eff}(\theta) \leq 15.0, \forall \theta$			
Θ_p (rad)	p region center colatitude	$\cos(\Theta_p) \sim U(-1, 1)$	$1.13_{-0.47}^{+0.63}$	0.13	0.425
Θ_s (rad)	s region center colatitude	$\cos(\Theta_s) \sim U(-1, 1)$	$1.98_{-0.62}^{+0.49}$	0.13	1.610
ϕ_p (cycles)	p region initial phase ^a	$\phi_p \sim U(-0.5, 0.5)$, wrapped	bimodal	3.46	-0.267
ϕ_s (cycles)	s region initial phase ^b	$\phi_s \sim U(-0.5, 0.5)$, wrapped	bimodal	3.47	-0.309
ζ_p (rad)	p region angular radius	$\zeta_p \sim U(0, \pi/2)$	$0.191_{-0.057}^{+0.102}$	1.69	0.233
ζ_s (rad)	s region angular radius	$\zeta_s \sim U(0, \pi/2)$	$0.189_{-0.058}^{+0.104}$	1.68	0.132
	no region-exchange degeneracy	$\Theta_s \geq \Theta_p$			
	non-overlapping hot regions	function of $(\Theta_p, \Theta_s, \phi_p, \phi_s, \zeta_p, \zeta_s)$			
$\log_{10}(T_p)$ [K]	p region NSX effective temperature	$\log_{10}(T_p) \sim U(5.1, 6.8)$, NSX limits	$6.014_{-0.063}^{+0.048}$	2.90	6.068
$\log_{10}(T_s)$ [K]	s region NSX effective temperature	$\log_{10}(T_s) \sim U(5.1, 6.8)$, NSX limits	$6.017_{-0.062}^{+0.048}$	2.89	6.086
$\cos(i)$	cosine Earth inclination to spin axis	$M, \cos(i) \sim N(\mu^*, \Sigma^*)$	$0.0426_{-0.0028}^{+0.0029}$	0.01	0.045
D (kpc)	Earth distance	$D \sim \text{skewnorm}(1.7, 1.0, 0.23)^c$	$1.19_{-0.14}^{+0.14}$	0.08	1.145
N_H [10^{20} cm^{-2}]	interstellar neutral H column density	$N_H \sim U(0, 10)$	$1.34_{-0.93}^{+1.87}$	1.07	0.029
α_{NICER}	NICER effective-area scaling	$\alpha_{NICER} \sim N(1, 0.15^2)$	$0.97_{-0.13}^{+0.15}$	0.04	1.083
	Sampling process information				
	number of free parameters: ^d 14				
	number of processes: ^e 1				
	number of live points: 4×10^4				
	hypervolume expansion factor: 0.1^{-1}				
	termination condition: 10^{-1}				
	number of core ^f hours: 24000				
	likelihood evaluations: 26858453				
	nested replacements: 1208371				
	effective sample size: 303905				

Notes. See the caption of Table 1 and the associated footnotes for details.^a With respect to the meridian on which Earth lies.^b With respect to the meridian on which the Earth antipode lies.^c The PDF defined as `scipy.stats.skewnorm.pdf(D, 1.7, loc = 1.002, scale = 0.227)`, truncated to the interval $D \in [0, 1.7]$ kpc.^d In the sampling space; the number of background count rate variables is equal to the number of channels defined by the NICER data set.^e The mode-separation MULTINEST variant was deactivated, meaning that isolated modes are not evolved independently and nested sampling threads contact multiple modes.^f Intel® Xeon® E5-2697A v4.

approximation error; a higher number of live points also yields finer sampling of the distribution. Potential variation arising from sampler choice is nicely illustrated in the exploratory study by Ashton et al. (2019), which investigates the effect on estimation of parameters of gravitational wave signals from compact binary coalescences.

These different modeling choices lead to some minor differences in the reported spreads of the radius posterior PDFs conditional on NICER data. However, the degree of overlap is high: we find $R = 11.29_{-0.81}^{+1.20}$ km (see Table 2); Miller et al. (2021) find $R = 11.51_{-1.13}^{+1.87}$ km.

The radius posterior PDF differences become more pronounced once the XMM data set is included. Our posterior is shifted down in radius relative to the Miller et al. (2021) posteriors which extend to higher radii. Once again, the two groups make a number of different choices that have more of an impact for a smaller data set. We use different formulations of the prior on the XMM background: Miller et al. (2021) use a

distribution based on the assumedly Poissonian observed numbers of blank-sky counts whereas we use a flat prior as described in Section 2.5.2. Our posterior is also conditional on a broader prior for the cross-calibration uncertainty of the two instruments than Miller et al. (who restrict the maximum relative calibration offset to $\pm 10\%$); this permits lower inferred radii in our analysis (see Section 4.2, where we study the effect of narrowing the cross-calibration uncertainty). And as already mentioned, Miller et al. are more agnostic in terms of priors on the radius (allowing $R > 16$ km): the XMM likelihood function permits larger radii (see Figure 7) and hence their posterior PDF extends accordingly to higher radii.

Despite these differences, the inclusion of the XMM data is still extremely valuable, because in both analyses there is a consistent increase in radius, with the lowest radii being ruled out. Moreover, there are good prospects for improving this: once the uncertainties on the NICER background estimates mentioned in Section 4.2.1 are clear, we anticipate being able

to use those to supplement the indirect constraint on the NICER background provided by the XMM data set.

5. Conclusions

We have derived a posterior distribution of the radius of the massive rotation-powered millisecond pulsar PSR J0740+6620, conditional on NICER XTI pulse-profile modeling, joint NANOGrav and CHIME/Pulsar wideband radio timing, and XMM EPIC spectroscopy. The radius that we infer for PSR J0740+6620 is $12.39^{+1.30}_{-0.98}$ km with an inferred mass (dominated by the radio-derived prior) of $2.072^{+0.067}_{-0.066} M_{\odot}$. A measurement of radius for such a high-mass pulsar should provide a strong constraint on dense matter EOS models (see Raaijmakers et al. 2021), with the derived radius favoring models of intermediate stiffness. We anticipate being able to improve this measurement in the near future thanks to the ongoing development of detailed models of the NICER background. This will be incorporated into future pulse-profile modeling, improving upon the current indirect constraint provided by the XMM data set.

Pulse-profile modeling also enables us to infer the properties of the X-ray-emitting hot regions, which are assumed to be linked to the magnetic field structure. The two hot regions are not antipodal, arguing against a simple dipole magnetic field. There is, however, no evidence for extended crescents, as indicated by pulse profile modeling for PSR J0030+0451 (Riley et al. 2019); simple circular hot regions (spherical caps) suffice to describe the PSR J0740+6620 data adequately. How this relates to the evolutionary history of the two sources remains to be determined.

The analysis presented here also includes improvements to our pulse-profile modeling methodology and software, most notably the ability to include (in this case) phase-averaged X-ray data from XMM EPIC. For PSR J0740+6620, the inclusion of this data set led to a remarkable shift in the inferred radius. We have also investigated the sensitivity to uncertainties in instrumental cross-calibration, an area where we may be able to improve our modeling in the future. XMM EPIC data sets exist for other NICER pulse-profile modeling targets, including the source analyzed in Riley et al. (2019), PSR J0030+0451. In that analysis the XMM EPIC data was used retrospectively, as a check on the consistency of the inferred model a posteriori; more formally, this information should be used to form a likelihood function that is a product of likelihood function factors over telescopes, as in this present work. In future work, we will use the improved pipeline presented in this Letter to perform joint analysis of the NICER and XMM data sets for PSR J0030+0451 and other sources.


















This work was supported in part by NASA through the NICER mission and the Astrophysics Explorers Program. T.E.R. and A.L.W. acknowledge support from ERC Consolidator grant No. 865768 AEONS (PI: Watts). T.E.R. also acknowledges support from the Nederlandse Organisatie voor Wetenschappelijk Onderzoek (NWO) through the VIDI and Projectruimte grants (PI: Nissanke). This work was sponsored by NWO Exact and Natural Sciences for the use of supercomputer facilities, and was carried out on the Dutch national e-infrastructure with the support of SURF Cooperative. This work was granted access to the HPC resources of CALMIP supercomputing center under the allocation 2016- P19056. S.B. was funded in part by NASA grants












NNX17AC28G and 80NSSC20K0275. S.M.M. thanks NSERC for support. 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. R.M.L. acknowledges the support of NASA through Hubble Fellowship Program grant HST-HF2-51440.001. 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. T.T.P. is a NANOGrav Physics Frontiers Center Postdoctoral Fellow funded by the National Science Foundation award number 1430284. The National Radio Astronomy Observatory is a facility of the National Science Foundation operated under cooperative agreement by Associated Universities, Inc. S.M.R. is a CIFAR Fellow and is supported by the NSF Physics Frontiers Center award 1430284. Pulsar research at UBC is supported by an NSERC Discovery Grant and by the Canadian Institute for Advanced Research. This research has made extensive use of NASA’s Astrophysics Data System Bibliographic Services (ADS) and the arXiv. We would also like to acknowledge the administrative and facilities staff whose labor supports our work.

Facility: NICER XTI (Gendreau et al. 2016), NANOGrav, Green Bank Telescope, CHIME/Pulsar, XMM-Newton EPIC.

Software: Python/C language (Oliphant 2007), GNU Scientific Library (GSL; Gough 2009), NumPy (van der Walt et al. 2011), Cython (Behnel et al. 2011), SciPy (Jones et al. 2001), OpenMP (Dagum & Menon 1998), MPI (Forum 1994), *MPI for Python* (Dalcín et al. 2008), Matplotlib (Hunter 2007; Droettboom et al. 2018), IPython (Perez & Granger 2007), Jupyter (Kluyver et al. 2016), TEMPO2 (photons; Hobbs et al. 2006), PINT (photonphase; <https://github.com/nanograv/PINT>), MULTINEST (Feroz et al. 2009), PYMULTINEST (Buchner et al. 2014), GetDist (Lewis 2019; <https://github.com/cmbant/getdist>), *nestcheck* (Higson 2018; Higson et al. 2018, 2019), *fgivenx* (Handley 2018), NICERsoft (<https://github.com/paulray/NICERsoft>), X-PSIv0.7 (<https://github.com/ThomasEdwardRiley/xpsi>; Riley 2021).

ORCID iDs

Thomas E. Riley  <https://orcid.org/0000-0001-9313-0493>
 Anna L. Watts  <https://orcid.org/0000-0002-1009-2354>
 Paul S. Ray  <https://orcid.org/0000-0002-5297-5278>
 Slavko Bogdanov  <https://orcid.org/0000-0002-9870-2742>
 Sebastien Guillot  <https://orcid.org/0000-0002-6449-106X>
 Sharon M. Morsink  <https://orcid.org/0000-0003-4357-0575>
 Anna V. Bilousova  <https://orcid.org/0000-0002-7177-6987>
 Devarshi Choudhury  <https://orcid.org/0000-0002-2651-5286>
 Julia S. Deneva  <https://orcid.org/0000-0003-1226-0793>
 Keith C. Gendreau  <https://orcid.org/0000-0001-7115-2819>
 Alice K. Harding  <https://orcid.org/0000-0001-6119-859X>
 Wynn C. G. Ho  <https://orcid.org/0000-0002-6089-6836>
 James M. Lattimer  <https://orcid.org/0000-0002-5907-4552>
 Renee M. Ludlam  <https://orcid.org/0000-0002-8961-939X>
 Craig B. Markwardt  <https://orcid.org/0000-0001-9803-3879>
 Takashi Okajima  <https://orcid.org/0000-0002-6054-3432>
 Chanda Prescod-Weinstein  <https://orcid.org/0000-0002-6742-4532>

Ronald A. Remillard  <https://orcid.org/0000-0003-4815-0481>
 Michael T. Wolff  <https://orcid.org/0000-0002-4013-5650>
 Emmanuel Fonseca  <https://orcid.org/0000-0001-8384-5049>
 H. Thankful Cromartie  <https://orcid.org/0000-0002-6039-692X>
 Matthew Kerr  <https://orcid.org/0000-0002-0893-4073>
 Timothy T. Pennucci  <https://orcid.org/0000-0001-5465-2889>
 Aditya Parthasarathy  <https://orcid.org/0000-0002-4140-5616>
 Scott Ransom  <https://orcid.org/0000-0001-5799-9714>
 Ingrid Stairs  <https://orcid.org/0000-0001-9784-8670>
 Lucas Guillemot  <https://orcid.org/0000-0002-9049-8716>
 Ismael Cognard  <https://orcid.org/0000-0002-1775-9692>

References

- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2018, *PhRvL*, **121**, 161101
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019, *PhRvX*, **9**, 011001
 Abbott, R., Abbott, T. D., Abraham, S., et al. 2020, *ApJL*, **896**, L44
 Akmal, A., Pandharipande, V. R., & Ravenhall, D. G. 1998, *PhRvC*, **58**, 1804
 AlGendy, M., & Morsink, S. M. 2014, *ApJ*, **791**, 78
 Al-Mamun, M., Steiner, A. W., Nättilä, J., et al. 2021, *PhRvL*, **126**, 061101
 Alvarez-Castillo, D., Ayriyan, A., Barnaföldi, G. G., & Pósfay, P. 2020, *PPN*, **51**, 725
 Annala, E., Gorda, T., Kurkela, A., Nättilä, J., & Vuorinen, A. 2020, *NatPh*, **16**, 907
 Antoniadis, J., Tauris, T. M., Özel, F., et al. 2016, arXiv:1605.01665
 Arzoumanian, Z., Brazier, A., Burke-Spolaor, S., et al. 2018, *ApJS*, **235**, 37
 Ashton, G., Hübner, M., Lasky, P. D., et al. 2019, *ApJS*, **241**, 27
 Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, **156**, 123
 Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, **558**, A33
 Badnell, N. R., Bautista, M. A., Butler, K., et al. 2005, *MNRAS*, **360**, 458
 Baillot d'Étivaux, N., Guillot, S., Margueron, J., et al. 2019, *ApJ*, **887**, 48
 Baym, G., Hatsuda, T., Kojo, T., et al. 2018, *RPPH*, **81**, 056902
 Behnel, S., Bradshaw, R., Citro, C., et al. 2011, *CSE*, **13**, 31
 Beronya, D. M., Karpova, A. V., Kirichenko, A. Y., et al. 2019, *MNRAS*, **485**, 3715
 Bilous, A. V., Watts, A. L., Harding, A. K., et al. 2019, *ApJL*, **887**, L23
 Biswas, B., Char, P., Nandi, R., & Bose, S. 2021, *PhRvD*, **103**, 103015
 Blaschke, D., Ayriyan, A., Alvarez-Castillo, D. E., & Grigorian, H. 2020, *Univ*, **6**, 81
 Bogdanov, S., Dittmann, A. J., Ho, W. C. G., et al. 2021, *ApJL*, **914**, L15
 Bogdanov, S., & Grindlay, J. E. 2009, *ApJ*, **703**, 1557
 Bogdanov, S., Guillot, S., Ray, P. S., et al. 2019a, *ApJL*, **887**, L25
 Bogdanov, S., Lamb, F. K., Mahmoodifar, S., et al. 2019b, *ApJL*, **887**, L26
 Brewer, B., & Foreman-Mackey, D. 2016, arXiv:1606.03757
 Buchner, J., Georgakakis, A., Nandra, K., et al. 2014, *A&A*, **564**, A125
 Chen, A. Y., Yuan, Y., & Vasilopoulos, G. 2020, *ApJL*, **893**, L38
 Christian, J.-E., & Schaffner-Bielich, J. 2020, *ApJL*, **894**, L8
 Clyde, M., Çetinkaya-Rundel, M., Rundel, C., et al. 2021, An Introduction to Bayesian Thinking: A Companion to the Statistics with R Course, <https://statswithr.github.io/book/>
 Colgan, J., Kilcrease, D. P., Magee, N. H., et al. 2016, *ApJ*, **817**, 116
 Cromartie, H. T., Fonseca, E., Ransom, S. M., et al. 2020, *NatAs*, **4**, 72
 Dagum, L., & Menon, R. 1998, *CSE*, **5**, 46
 Dalcín, L., Paz, R., Storti, M., & D'Elia, J. 2008, *JPDC*, **68**, 655
 De Grandis, D., Turolla, R., Wood, T. S., et al. 2020, *ApJ*, **903**, 40
 De, S., Finstad, D., Lattimer, J. M., et al. 2018, *PhRvL*, **121**, 091102
 Dexheimer, V., Gomes, R. O., Klähn, T., Han, S., & Salinas, M. 2021, *PhRvC*, **103**, 025808
 Dietrich, T., Coughlin, M. W., Pang, P. T. H., et al. 2020, *Sci*, **370**, 1450
 Drischler, C., Han, S., Lattimer, J. M., et al. 2021, *PhRvC*, **103**, 045808
 Droettboom, M., Caswell, T. A., Hunter, J., et al. 2018, matplotlib/matplotlib v2.2.2, Zenodo, doi:10.5281/zenodo.1202077
 Essick, R., Tews, I., Landry, P., Reddy, S., & Holz, D. E. 2020, *PhRvC*, **102**, 055803
 Fattoyev, F. J., Horowitz, C. J., Piekarewicz, J., & Reed, B. 2020, *PhRvC*, **102**, 065805
 Feroz, F., & Hobson, M. P. 2008, *MNRAS*, **384**, 449
 Feroz, F., Hobson, M. P., & Bridges, M. 2009, *MNRAS*, **398**, 1601
 Feroz, F., Hobson, M. P., Cameron, E., & Pettitt, A. N. 2013, arXiv:1306.2144
 Foight, D. R., Güver, T., Özel, F., & Slane, P. O. 2016, *ApJ*, **826**, 66
 Fonseca, E., Cromartie, H. T., Pennucci, T. T., et al. 2021a, *ApJL*, **915**, L12
 Fonseca, E., Cromartie, H. T., Pennucci, T. T., et al. 2021b, Refined Mass and Geometric Measurements of the High-Mass PSR J0740+6620: Probability Density Functions and their Credible Intervals, Zenodo, doi:10.5281/zenodo.4773599
 Forum, M. P. 1994, MPI: A Message-Passing Interface Standard, Knoxville, TN, <https://www.mpi-forum.org/docs/mpi-1.0/mpi-10.ps>
 Gelman, A., Carlin, J., Stern, H., et al. 2013, Bayesian Data Analysis, Third Edition, Chapman and Hall/CRC Texts in Statistical Science (London: Taylor and Francis)
 Gendreau, K. 2020, NICER Background Estimator Tools, https://heasarc.gsfc.nasa.gov/docs/nicer/tools/nicer_bkg_est_tools.html
 Gendreau, K. C., Arzoumanian, Z., Adkins, P. W., et al. 2016, *Proc. SPIE*, **9905**, 99051H
 Godzieba, D. A., Radice, D., & Bernuzzi, S. 2021, *ApJ*, **908**, 122
 González-Caniulef, D., Guillot, S., & Reisenegger, A. 2019, *MNRAS*, **490**, 5848
 Gough, B. 2009, GNU Scientific Library Reference Manual (3rd ed.; Surrey: Network Theory Ltd.)
 Greif, S. K., Raaijmakers, G., Hebel, K., Schwenk, A., & Watts, A. L. 2019, *MNRAS*, **485**, 5363
 Han, S., & Prakash, M. 2020, *ApJ*, **899**, 164
 Handley, W. 2018, *JOSS*, **3**, 849
 He, C., Ng, C. Y., & Kaspi, V. M. 2013, *ApJ*, **768**, 64
 Hebel, K. 2021, *PhR*, **890**, 1
 Hebel, K., Lattimer, J. M., Pethick, C. J., & Schwenk, A. 2013, *ApJ*, **773**, 11
 Higson, E. 2018, *JOSS*, **3**, 916
 Higson, E., Handley, W., Hobson, M., & Lasenby, A. 2018, *BayAn*, **13**, 873
 Higson, E., Handley, W., Hobson, M., & Lasenby, A. 2019, *MNRAS*, **483**, 2044
 HI4PI Collaboration, Ben Bekhti, N., Flöer, L., et al. 2016, *A&A*, **594**, A116
 Ho, W. C. G., & Heinke, C. O. 2009, *Natur*, **462**, 71
 Ho, W. C. G., & Lai, D. 2001, *MNRAS*, **327**, 1081
 Hobbs, G. B., Edwards, R. T., & Manchester, R. N. 2006, *MNRAS*, **369**, 655
 Hogg, D. W. 2012, arXiv:1205.4446
 Hogg, D. W., Price-Whelan, A. M., & Leistedt, B. 2020, arXiv:2005.14199
 Huang, K., Hu, J., Zhang, Y., & Shen, H. 2020, *ApJ*, **904**, 39
 Hunter, J. D. 2007, *CSE*, **9**, 90
 Iglesias, C. A., & Rogers, F. J. 1996, *ApJ*, **464**, 943
 Igoshchev, A., Verbunt, F., & Cator, E. 2016, *A&A*, **591**, A123
 Ishida, M., Tsujimoto, M., Kohmura, T., et al. 2011, *PASJ*, **63**, S657
 Jiang, J.-L., Tang, S.-P., Wang, Y.-Z., Fan, Y.-Z., & Wei, D.-M. 2020, *ApJ*, **892**, 55
 Jones, E., Oliphant, T., Peterson, P., et al. 2001, SciPy: Open Source Scientific Tools for Python, <http://www.scipy.org/>
 Kalapotharakos, C., Wadiasingh, Z., Harding, A. K., & Kazanas, D. 2021, *ApJ*, **907**, 63
 Kluyver, T., Ragan-Kelley, B., Pérez, F., et al. 2016, in Positioning and Power in Academic Publishing: Players, Agents and Agendas, ed. F. Loizides & B. Schmidt (Amsterdam: IOS Press), 87
 Kondratyev, I. A., Moiseenko, S. G., Bisnovaty-Kogan, G. S., & Glushikhina, M. V. 2020, *MNRAS*, **497**, 2883
 Kullback, S., & Leibler, R. A. 1951, *Ann. Math. Statist.*, **22**, 79
 Lalletment, R., Capitanio, L., Ruiz-Dern, L., et al. 2018, *A&A*, **616**, A132
 Landry, P., Essick, R., & Chatziioannou, K. 2020, *PhRvD*, **101**, 123007
 Lattimer, J. M., & Prakash, M. 2016, *PhR*, **621**, 127
 Lewis, A. 2019, arXiv:1910.13970
 Li, A., Jiang, J.-L., Tang, S.-P., et al. 2020, arXiv:2009.12571
 Lim, Y., Bhattacharya, A., Holt, J. W., & Pati, D. 2020, arXiv:2007.06526
 Lommen, A. N., Zepka, A., Backer, D. C., et al. 2000, *ApJ*, **545**, 1007
 Lorimer, D. R., & Kramer, M. 2004, Handbook of Pulsar Astronomy, Vol. 4 (Cambridge: Cambridge Univ. Press)
 Ludlam, R. M., Miller, J. M., Bachetti, M., et al. 2017, *ApJ*, **836**, 140
 Madsen, K. K., Beardmore, A. P., Forster, K., et al. 2017, *AJ*, **153**, 2
 Margalit, B., & Metzger, B. D. 2017, *ApJL*, **850**, L19
 Martino, L., Elvira, V., & Louzada, F. 2017, *SigPr*, **131**, 386
 Maselli, A., Sabatucci, A., & Benhar, O. 2021, *PhRvC*, **103**, 065804
 McLerran, L., & Reddy, S. 2019, *PhRvL*, **122**, 122701
 Miller, M. C. 2016, *ApJ*, **822**, 27
 Miller, M. C., Lamb, F. K., Dittmann, A. J., et al. 2019, *ApJL*, **887**, L24
 Miller, M. C., Lamb, F. K., Dittmann, A. J., et al. 2021, *ApJL*, in press (doi:10.3847/2041-8213/ac089b)

- 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
- Oertel, M., Hempel, M., Klähn, T., & Typel, S. 2017, *RvMP*, **89**, 015007
- Oliphant, T. E. 2007, *CSE*, **9**, 10
- Özel, F., & Freire, P. 2016, *ARA&A*, **54**, 401
- Pavlov, G. G., & Zavlin, V. E. 1997, *ApJL*, **490**, L91
- Perez, F., & Granger, B. E. 2007, *CSE*, **9**, 21
- Plucinsky, P. P., Beardmore, A. P., Foster, A., et al. 2017, *A&A*, **597**, A35
- Raaijmakers, G., Greif, S. K., Hebeler, K., et al. 2021, *ApJL*, in press (doi:10.3847/2041-8213/ac089a)
- Raaijmakers, G., Greif, S. K., Riley, T. E., et al. 2020, *ApJL*, **893**, L21
- Raaijmakers, G., Riley, T. E., Watts, A. L., et al. 2019, *ApJL*, **887**, L22
- Ray, P. S., Arzoumanian, Z., Ballantyne, D., et al. 2019, arXiv:1903.03035
- 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
- Riley, T. E. 2019, PhD Thesis, Univ. Amsterdam, <https://hdl.handle.net/11245.1/aa86fcf3-2437-4bc2-810e-cf9f30a98f7a>
- Riley, T. E. 2021, X-PSI: X-ray Pulse Simulation and Inference, v0.7, Astrophysics Source Code Library, ascl:2102.005
- 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, A NICER View of the Massive Pulsar PSR J0740+6620 Informed by Radio Timing and XMM-Newton Spectroscopy: Nested Samples for Millisecond Pulsar Parameter Estimation, v1.0.0, Zenodo, doi:10.5281/zenodo.4697625
- Ruiz, M., Shapiro, S. L., & Tsokaros, A. 2018, *PhRvD*, **97**, 021501
- Schertler, K., Greiner, C., Schaffner-Bielich, J., & Thoma2, M. H. 2000, *NuPhA*, **677**, 463
- Sedrakian, A., Weber, F., & Li, J. J. 2020, *PhRvD*, **102**, 041301
- Shibata, M., Zhou, E., Kiuchi, K., & Fujibayashi, S. 2019, *PhRvD*, **100**, 023015
- Shklovskii, I. S. 1970, *SvA*, **13**, 562
- Silva, H. O., Pappas, G., Yunes, N., & Yagi, K. 2021, *PhRvD*, **103**, 063038
- Sivia, D. S., & Skilling, J. 2006, *Data Analysis—A Bayesian Tutorial* (2nd ed.; Oxford: Oxford Univ. Press)
- Steiner, A. W., Heinke, C. O., Bogdanov, S., et al. 2018, *MNRAS*, **476**, 421
- Stergioulas, N., & Friedman, J. L. 1995, *ApJ*, **444**, 306
- Suvorov, A. G., & Melatos, A. 2020, *MNRAS*, **499**, 3243
- Tan, H., Noronha-Hostler, J., & Yunes, N. 2020, *PhRvL*, **125**, 261104
- Tang, S.-P., Jiang, J.-L., Gao, W.-H., Fan, Y.-Z., & Wei, D.-M. 2021, *PhRvD*, **103**, 063026
- Tews, I., Pang, P. T. H., Dietrich, T., et al. 2021, *ApJL*, **908**, L1
- Tolos, L., & Fabbietti, L. 2020, *PrPNP*, **112**, 103770
- Traversi, S., Char, P., & Pagliara, G. 2020, *ApJ*, **897**, 165
- Trotta, R. 2008, *ConPh*, **49**, 71
- Tsokaros, A., Ruiz, M., & Shapiro, S. L. 2020, *ApJ*, **905**, 48
- van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, *CSE*, **13**, 22
- Watts, A. L. 2019, in AIP Conf. Proc. 2127, Xiamen-CUSTIPEN Workshop on the EOS of Dense Neutron-Rich Matter in the Era of Gravitational Wave Astronomy, ed. A. Li, B.-A. Li, & F. Xu (Melville, NY: AIP), 020008
- Watts, A. L., Andersson, N., Chakrabarty, D., et al. 2016, *RvMP*, **88**, 021001
- Watts, A. L., Yu, W., Poutanen, J., et al. 2019, *SCPMA*, **62**, 29503
- Wilms, J., Allen, A., & McCray, R. 2000, *ApJ*, **542**, 914
- Wolff, M. T., Guillot, S., Bogdanov, S., et al. 2021, *ApJL*, in press (doi:10.3847/2041-8213/ac158e)
- Xie, W.-J., & Li, B.-A. 2020, *ApJ*, **899**, 4
- Xie, W.-J., & Li, B.-A. 2021, *PhRvC*, **103**, 035802
- Yang, J., & Piekarewicz, J. 2020, *ARNPS*, **70**, 21
- Zhang, N.-B., & Li, B.-A. 2020a, *ApJ*, **893**, 61
- Zhang, N.-B., & Li, B.-A. 2020b, *ApJ*, **902**, 38
- Zhang, S., Santangelo, A., Feroci, M., et al. 2019, *SCPMA*, **62**, 29502
- Zimmerman, J., Carson, Z., Schumacher, K., Steiner, A. W., & Yagi, K. 2020, arXiv:2002.03210