



PSR J0030+0451 Mass and Radius from *NICER* Data and Implications for the Properties of Neutron Star Matter

M. C. Miller¹ , F. K. Lamb^{2,3} , A. J. Dittmann¹ , S. Bogdanov⁴ , Z. Arzoumanian⁵, K. C. Gendreau⁵, S. Guillot^{6,7} , A. K. Harding⁸ , W. C. G. Ho^{9,10} , J. M. Lattimer¹¹ , R. M. Ludlam^{12,19} , S. Mahmoodifar^{1,8} , S. M. Morsink¹³ , P. S. Ray¹⁴ , T. E. Strohmayer¹⁵ , K. S. Wood¹⁶ , T. Enoto¹⁷ , R. Foster¹⁸ , T. Okajima⁵, G. Prigozhin¹⁸, and Y. Soong⁵

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

²Center for Theoretical Astrophysics and Department of Physics, University of Illinois at Urbana-Champaign, 1110 West Green Street, Urbana, IL 61801-3080, USA

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

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

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

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

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

⁸Astrophysics Science Division, NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA

⁹Department of Physics and Astronomy, Haverford College, 370 Lancaster Avenue, Haverford, PA 19041, USA

¹⁰Mathematical Sciences, Physics and Astronomy, and STAG Research Centre, University of Southampton, Southampton SO17 1BJ, UK

¹¹Department of Physics and Astronomy, Stony Brook University, Stony Brook, NY 11794-3800, USA

¹²Cahill Center for Astronomy and Astrophysics, California Institute of Technology, Pasadena, CA 91125, USA

¹³Department of Physics, University of Alberta, 4-183 CCIS, Edmonton, AB T6G 2E1, Canada

¹⁴Space Science Division, U.S. Naval Research Laboratory, Washington, DC 20375, USA

¹⁵Astrophysics Science Division and Joint Space-Science Institute, NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA

¹⁶Praxis, resident at the U.S. Naval Research Laboratory, Washington, DC 20375, USA

¹⁷The Hakubi Center for Advanced Research and Department of Astronomy, Kyoto University, Kyoto 606-8302, Japan

¹⁸MIT Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, 70 Vassar Street, Cambridge, MA, USA

Received 2019 August 9; revised 2019 October 9; accepted 2019 October 24; published 2019 December 12

Abstract

Neutron stars are not only of astrophysical interest, but are also of great interest to nuclear physicists because their attributes can be used to determine the properties of the dense matter in their cores. One of the most informative approaches for determining the equation of state (EoS) of this dense matter is to measure both a star's equatorial circumferential radius R_e and its gravitational mass M . Here we report estimates of the mass and radius of the isolated 205.53 Hz millisecond pulsar PSR J0030+0451 obtained using a Bayesian inference approach to analyze its energy-dependent thermal X-ray waveform, which was observed using the *Neutron Star Interior Composition Explorer* (*NICER*). This approach is thought to be less subject to systematic errors than other approaches for estimating neutron star radii. We explored a variety of emission patterns on the stellar surface. Our best-fit model has three oval, uniform-temperature emitting spots and provides an excellent description of the pulse waveform observed using *NICER*. The radius and mass estimates given by this model are $R_e = 13.02_{-1.06}^{+1.24}$ km and $M = 1.44_{-0.14}^{+0.15} M_\odot$ (68%). The independent analysis reported in the companion paper by Riley et al. explores different emitting spot models, but finds spot shapes and locations and estimates of R_e and M that are consistent with those found in this work. We show that our measurements of R_e and M for PSR J0030+0451 improve the astrophysical constraints on the EoS of cold, catalyzed matter above nuclear saturation density.

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

1. Introduction

A key current goal of nuclear physics is to understand the properties of cold catalyzed matter above the saturation density of nuclear matter. Matter at these densities cannot be studied in terrestrial laboratories. Hence observations of neutron stars—which contain large quantities of such matter—play a key role (see, e.g., Lattimer & Prakash 2007). Over the last few years, the discovery of several high-mass neutron stars (Demorest et al. 2010; Antoniadis et al. 2013; Arzoumanian et al. 2018a; Cromartie et al. 2019) and measurement of the binary tidal deformability during a neutron star merger (Abbott et al. 2017, 2018; De et al. 2018) have advanced our knowledge of the properties of cold dense matter, but precise and reliable measurements of neutron star radii would significantly improve our understanding.

Various radius estimates have been made using models of the X-ray emission from quiescent neutron stars (see Steiner et al. 2018 for a recent summary), from neutron stars during thermonuclear X-ray bursts (see Steiner et al. 2010; Özel et al. 2016; and Nättilä et al. 2017 for different perspectives), and from accretion-powered millisecond pulsars (see Salmi et al. 2018), with inferred radii typically ranging from ~ 10 km to ~ 14 km, consistent with most theoretical predictions. However, these estimates are susceptible to significant systematic errors, in the sense that a model could provide a formally good fit to the data but yield a credible interval for the radius that strongly excludes the true value (Miller 2013; Miller & Lamb 2016).

In contrast, analyses of the soft X-ray pulse waveforms observed using the *Neutron Star Interior Composition Explorer* (*NICER*) are expected to be less susceptible to systematic errors. Analyses of synthetic waveforms carried out prior to the

¹⁹ Einstein Fellow.

launch of *NICER* showed that, for the cases considered, using model assumptions different from the true situation (e.g., different emission or beaming patterns, different spectra, or different surface temperature distributions) did not significantly bias parameter estimates, provided that the fit was formally good (Lo et al. 2013; Miller & Lamb 2015). Simple pulse waveform models have previously been fit to the soft X-ray waveforms of rotation-powered pulsars observed using *ROSAT* and the *Extreme Ultraviolet Explorer* (*EUVE*; Pavlov & Zavlin 1997; Zavlin & Pavlov 1998) and *XMM-Newton* (see, e.g., Bogdanov et al. 2007, 2008; Bogdanov 2013). These fits gave estimates for the radii of these pulsars that were consistent with the expected range of neutron-star radii, but the number of counts available was too small to obtain tight constraints. Nevertheless, these pioneering studies indicated that rotation-powered pulsars are promising targets for such measurements.

Here we present the results of an analysis by *NICER* team members of the currently available data on the 205.53 Hz (Arzoumanian et al. 2018a) millisecond pulsar PSR J0030+0451. The results of an independent analysis of these data by other *NICER* team members are presented in the companion paper by Riley et al. (2019). The analysis presented there uses an independently developed code for computing pulse waveforms, a different approach to modeling the heated regions on the stellar surface, and different sampling methods than the ones used in the analysis presented here. The consistency of the surface temperature distributions and mass–radius posteriors reported there with those reported here suggests that the unavoidable uncertainties in models of the heated areas on the stellar surface have not greatly affected the results of either analysis. In Section 2 we describe how we obtained and processed the *NICER* data on PSR J0030+0451. In Section 3 we discuss our methods, including our modeling of the emission from the stellar surface, our approach to modeling the soft X-ray waveform, our parameter estimation and model evaluation methods, and our treatment of instrumental effects and bias. In Section 4 we describe how we used analyses of synthetic waveforms to verify our parameter estimation codes and procedures, and to evaluate the consequences of our decision to analyze a subset of the *NICER* data on PSR J0030+0451. In Section 5 we present and discuss our best estimates of the radius and mass of PSR J0030+0451 using the currently available *NICER* data on its pulse waveform. In Section 6 we compute and discuss the constraints on two different parameterized models of the equation of state (EoS) of neutron star matter implied by our measurements of the radius and mass of PSR J0030+0451. We summarize our conclusions in Section 7.

2. Observations

2.1. Data Used

Here, we briefly summarize the *NICER* X-ray Timing Instrument (XTI) data set that was used for the parameter estimation analyses presented here. More detailed descriptions of the observations, data reduction, and event folding procedures, and a thorough examination of the properties of the final event data can be found in Bogdanov et al. (2019a).

The *NICER* XTI data on PSR J0030+0451 were collected in a set of short observations (typically lasting a few hundred to ~ 2000 s) over the period from 2017 July 24 to 2018 December 9. Due to the exceptional rotational stability of PSR J0030

+0451 and the superb absolute timing accuracy of *NICER*, neither the large time span nor the discontinuities in the exposure had any adverse effects, such as smearing due to long-term pulse-phase drifts, on the pulse profile.

2.2. Filtering of the Data and Event Folding

For this analysis, we used products from the calibration database (CALDB) version 20181105 and gain solution version `optmv7`. In addition to the standard pipeline processing event filtering that all *NICER* data undergo, the PSR J0030+0451 data set was subjected to screening criteria tailored to the inference analyses discussed later.

We included only data accumulated when the source subtended an angle greater than 80° from the Sun, in order to minimize the contamination from optical loading in the lowest-energy detector channels. All events from the XTI detector with DET_ID 34 were excised because it frequently exhibits significantly elevated count rates relative to the other detectors (analogous to a “hot pixel” in charge-coupled device detectors).

In the final filtering stage, we constructed a time series with 16 s time bins using the event data in channels 25–800 and then discarded all the events in those time bins that had an average count rate greater than 3 counts s^{-1} . The long-term average count rate during the *NICER* observations of PSR J0030+0451 was $\sim 0.7 \text{ counts s}^{-1}$. This cut therefore excluded the event data in all the 16 s time bins that had count rates more than four times the long-term average count rate. The purpose of this cut was to improve the signal-to-noise ratio of the measured thermal emission from the pulsar by removing data contaminated by flares that were not filtered out by the standard pipeline processing.

The thermal emission pulse profile that was used to estimate the mass and radius of PSR J0030+0451 was constructed by averaging the data from the roughly 400 million rotation periods observed by *NICER* during more than 18 months of observations. We compared the profile measured during the first third of this observing time with the profile measured during the second third, and also compared the profile measured during the first half of this observing time with the profile measured during the second half. These comparisons showed no statistically significant differences between the compared profiles. We expect that any shorter-term flux variations were averaged out by using such a long data set.

To generate the final X-ray pulse profile as a function of pulsar rotational phase and energy, we used the “photons” plugin for the Tempo2 pulsar timing package (Hobbs et al. 2006) and the best publicly available radio timing solution, obtained by NANOGrav (Arzoumanian et al. 2018b), to assign pulse phases to each event. The number of the “pulse invariant” (PI) detector channel provides an estimate of the photon energy: the photon energy in eV is approximately equal to $10 \times$ the PI channel number in which the event occurs.

As discussed in Bogdanov et al. (2019a), the final X-ray pulse profile was cross-verified by comparing it with an independently generated profile in which pulse phases were assigned to events using the PINT pulsar timing software. The pulse phases assigned to individual events by the two procedures differed by $\lesssim 1 \mu\text{s}$, and the differences between the two folded X-ray pulse profiles were negligible. The final event list used in the analysis reported in this Letter is available at <https://doi.org/10.5281/zenodo.3524457>.

2.3. Calibration of the Instrument

To incorporate the performance of the *NICER* XTI in our model, we used version 1.02 of the redistribution matrix file (RMF), which gives the probability that a photon in a particular energy range is registered by a particular detector channel, and version 1.02 of the ancillary response file (ARF), which provides the effective area of the telescope as a function of energy, as measured on axis.

Spectral calibration of the *NICER* XTI has been carried out primarily using observations of the Crab pulsar and nebula. The response below channel 40 is more uncertain than the response at higher energies, because of current imperfections in the modeling of instrumental effects such as the detector trigger efficiency. Also, because at these low energies the Crab spectrum is strongly attenuated by the interstellar medium, only a relatively small number of counts are available to determine the spectrum at these energies. Fits to *NICER* observations of the Crab pulsar and nebula in the energy channels 40–299 used in this study (see Section 5.2) show systematic residuals, but these are typically $\lesssim 5\%$ (Ludlam et al. 2018). They are likely due to imperfect modeling of the microphysics of the concentrator optics and astrophysical features, such as oxygen edges and emission lines.

3. Methods

In this section, we first describe our modeling of the soft X-ray emission from hot spots on the stellar surface and our approach to modeling the soft X-ray pulses produced by the rotation of these spots. We then present the specific pulse waveform models that we consider in this work and our procedure for constructing and fitting these model waveforms to the *NICER* pulse waveform data on PSR J0030+0451. We end this section by discussing how we deal with the errors in the response of the *NICER* instrument and the evidence that including data from channels below channel 40 would risk biasing our estimates of the mass and radius of PSR J0030+0451. We defer to subsequent sections discussion of our use of synthetic pulse waveform data to verify our parameter estimation procedures, our use of *NICER* pulse waveform data to estimate the mass and radius of PSR J0030+0451, our assessment of the reliability of these estimates, and their implications for the EoS of cold dense matter.

3.1. Modeling the Emission from the Stellar Surface

The soft X-ray pulse waveform of PSR J0030+0451 is thought to be produced by hot regions on its rotating surface created by inward moving particles created in electron-positron pair cascades interacting with the outermost layers of the neutron star (Arons 1981; Harding & Muslimov 2001). Models of global pulsar magnetospheres (see, e.g., Kalapotharakos et al. 2014; Philippov et al. 2015; Brambilla et al. 2018) have delineated the current-closure patterns of pulsars, which show outward currents in the magnetic polar regions balanced by return currents over the outer parts of the polar caps that are connected to the current sheet. Some of the particles that make up the return currents (those that are accelerated in the current sheet) are responsible for the pulsar’s γ -ray and nonthermal X-ray emission (Kalapotharakos et al. 2018; Philippov & Spitkovsky 2018). The polar cap pair cascades that take place in both outgoing and return current regions (Timokhin &

Arons 2013) supply pairs for radio emission. Some of the particles produced in these pair cascades are accelerated downward and are expected to deposit energy deep in the neutron star’s atmosphere (see Tsai 1974 and Bogdanov et al. 2007). This energy is expected to heat the atmosphere at large optical depths, producing thermal soft X-ray emission from the stellar surface (Harding & Muslimov 2001, 2002). The spectrum of this emission is expected to depend on the chemical composition of the star’s atmosphere and, possibly, on the strength and orientation of the magnetic field at the locations of the hot spots (Potekhin 2014).

The models of the thermal emission from these hot spots that we have used in the analysis of the soft X-ray waveform of PSR J0030+0451 that we report here assume that the upper atmosphere of the pulsar is fully ionized pure hydrogen and that, for our purposes, the magnetic fields in the hot spots can be neglected. We now discuss these assumptions.

Composition of the upper atmosphere. Our model atmospheres were constructed assuming that the gas in the atmosphere is fully ionized, and hence that the dominant opacity is free–free absorption (Ho & Lai 2001). While opacity tables for partially ionized matter exist (Iglesias & Rogers 1996; Badnell et al. 2005; Colgan et al. 2016), they do not cover the full range of temperatures and energies needed for the present analysis. However, at the temperatures $\sim 10^6$ K that are relevant for the present study, the fraction of hydrogen that is neutral is very low; comparison of the outward specific intensity given by our fully ionized hydrogen atmosphere model differs by at most 1%–2% at 0.5–1 keV from that for an atmosphere with an effective temperature of 10^6 K given by a model constructed using the Opacity Project (OP) opacity table (Badnell et al. 2005), which allows for partial ionization (see Bogdanov et al. 2019b).

Our assumption that the upper atmosphere is pure hydrogen is based primarily on the astrophysical abundance of hydrogen and calculations that show how the lightest element present in a neutron star atmosphere will float to the top within seconds. These calculations are extrapolations of earlier calculations by, for example, Alcock & Illarionov (1980). Given the astrophysical abundance of hydrogen and its tendency to float to the top, we expect it to completely dominate the composition of the upper atmosphere.

If the outer atmosphere of the pulsar were instead dominated by helium, the emission spectrum and beaming pattern in the energy range 0.5–1.0 keV would be expected to be very similar to those for a hydrogen atmosphere, but their normalizations would be slightly smaller than for hydrogen (see Bogdanov et al. 2019b). In particular, hydrogen and helium atmospheres with an effective temperature of 10^6 K (the approximate inferred temperature of the heated regions on the surface of PSR J0030+0451) are predicted to produce specific intensities normal to the surface that differ from each other by $\lesssim 5\%$ in the 0.5–1 keV energy range (again see Bogdanov et al. 2019b). Our current best fit of the zero-field NSX helium atmosphere (Ho & Lai 2001) to the *NICER* data on PSR J0030+0451 appears disfavored relative to our current best fit of the zero-field NSX hydrogen atmosphere (Ho & Lai 2001; see also Ho & Heinke 2009). Specifically, our current fit using the NSX helium model yields a best-fit mass of $2.7 M_{\odot}$, with little posterior probability below $2.0 M_{\odot}$. Although more work would be needed to completely rule out a helium atmosphere

for PSR J0030+0451, this mass estimate is so large that we interpret it as disfavoring a helium atmosphere.

Thus, while there is at present no direct spectroscopic confirmation that the upper atmosphere of PSR J0030+0451 is pure hydrogen, the fact that the spectrum and beaming pattern predicted by hydrogen atmosphere models agree well with the spectrum and beaming of the soft X-ray emission from PSR J0030+0451 observed using *NICER*, and that the X-ray spectrum and beaming predicted by helium atmosphere models appear disfavored, supports our assumption that the upper atmosphere is hydrogen. Therefore, in computing model waveforms we used the local emission for a given surface gravity, effective temperature, angle from the surface normal, and photon energy that is predicted by the NSX code developed by Ho & Lai (2001; see also Ho & Heinke 2009) for a fully ionized hydrogen atmosphere.

The 205.53 Hz rotation frequency of PSR J0030+0451 (Arzoumanian et al. 2018a) is small enough that the only effects of the rotation that need to be taken into account in computing pulse waveforms are the special relativistic Doppler shift and aberration caused by the motion of the emitting surface and the rotationally induced oblateness of the star; the effect of frame-dragging and the effect on the exterior spacetime of the rotationally induced stellar mass quadrupole are negligible (see Morsink et al. 2007 and Bogdanov et al. 2019b). We therefore used the so-called oblate-star Schwarzschild-spacetime (OS) approximation (Cadeau et al. 2007; Morsink et al. 2007; AIGendy & Morsink 2014) when tracing light rays from the stellar surface to the observer. Computations of pulse waveforms like those used here were verified by members of the *NICER* Lightcurve Working Group, as described in Bogdanov et al. (2019b).

Strength and structure of the stellar magnetic field. We can estimate an approximate lower bound on the strength of the surface magnetic field of PSR J0030+0451 by assuming that the field is a centered dipole and estimating the strength of the dipole moment using the best braking models currently available (see, e.g., Contopoulos & Spitkovsky 2006). The period P of PSR J0030+0451 is 0.00487 seconds and its period derivative \dot{P} is 1.02×10^{-20} (Arzoumanian et al. 2018a). If we use Equation (12) of Contopoulos & Spitkovsky (2006) and, for the sake of argument, choose a magnetic field inclination θ of $\pi/2$ and the field configuration described by $\alpha = 0$, then the inferred strength of the magnetic field at the stellar surface is $B \sim 4 \times 10^8$ G. In principle, this represents an approximate lower bound on the strength of the total magnetic field in the atmospheres of the hot spots. For comparison, the electron cyclotron energy in a magnetic field of this strength is $\hbar\omega_c \approx 4$ eV, which is more than 50 times smaller than the lowest photon energies that can be observed using *NICER* and ~ 100 times smaller than the lowest energies that we consider in our fits.

We note, however, that the three hot spots in the “southern” rotational hemisphere (we denote the rotational hemisphere that includes the line of sight to the observer as the “northern” rotational hemisphere) that can explain the pulse waveform of PSR J0030+0451 observed using *NICER* suggest that the magnetic field of PSR J0030+0451 is at least an offset dipole, and may have higher moments that dominate the field strength near the stellar surface. A complex field would be consistent with the structures of the magnetic fields of other millisecond rotation-powered pulsars inferred from their radio and γ -ray

emission (Ruderman 1991; Ruderman et al. 1998), and the inferred magnetic field configurations of the accretion-powered millisecond pulsars that are thought to be the progenitors of the millisecond rotation-powered pulsars (Lamb et al. 2009a, 2009b).

If the magnetic field of PSR J0030+0451 does have significant higher multipole moments, they may not be negligible at the light cylinder, which is closer to the stellar surface in millisecond pulsars than in pulsars with lower spin rates. Moreover, in some spin-down models (again see Contopoulos & Spitkovsky 2006), the spin-down rate depends on the magnetic flux enclosed by the open-field-line region, which is defined by the boundary of the closed magnetosphere and may not extend all the way out to the light cylinder. The way this flux evolves with spin rate can be affected by the presence of multipolar field components, complicating the relation between the strength of the surface magnetic field and the spin-down rate.

Given this complexity, it is difficult to infer the strength and configuration of the magnetic field on the surface of PSR J0030+0451. In current braking models, the braking torque scales with the strength of the dipole component of the magnetic field at some characteristic distance from the star. For a given strength of the dipole component at this distance, the strength of the magnetic field at the stellar surface scales as $(R_e/d)^3$, where R_e is the circumferential radius of the star and $2d$ is the distance through the star from one magnetic pole to the other. Thus, if the large-scale magnetic field is essentially dipolar and the corresponding magnetic poles are a significant distance south of the rotational equator, the field strength at the surface of the star could be 5–10 \times greater than it would be for a centered dipole field with the same strength a large distance from the stellar surface. In this case, the indicated strength of the magnetic fields at the hot spots could be as large as $B \sim 4 \times 10^9$ G.

Effect of the magnetic field on the soft X-ray emission from the stellar surface. The model atmospheres that we have used to compute the waveforms that we fitted to the *NICER* waveform data assume that the effects of the stellar magnetic field on the structure and radiative properties of the atmosphere are negligible. This approximation is valid for magnetic field strengths that are less than $B_c = e^3 m_e^2 c / \hbar^3 = 2 \times 10^9$ G. The upper bound on the strength of the magnetic field that we estimated above by assuming that it is dipolar and using the best currently available braking models is an order of magnitude smaller than B_c , but our estimates of the possible field strengths at the hot spots are a factor of two larger than B_c .

If the magnetic field exceeds B_c , its effect is non-perturbative and dominant (see Lai 2001): the radiation field becomes polarized, the beaming pattern becomes strongly anisotropic, and spectral features associated with the electron cyclotron resonance and bound species become important (Potekhin 2014). The absence of any obvious spectral feature in the spectra observed using *XMM-Newton* (see, e.g., Bogdanov & Grindlay 2009) and in the ~ 0.3 –3 keV spectra obtained using *NICER* rules out $B = (3$ –30) $\times 10^{10}$ G, assuming a redshift of 20%. In addition, while an electron cyclotron line produced by a magnetic field $B \gg 10^{11}$ G could be hidden at $E > 3$ keV, such fields would produce pencil or fan beams that are unlikely to be able to produce an energy-dependent pulse profile that can match the profile observed using *NICER*.

The changes in the spectrum and the beaming pattern that occur for magnetic fields $B \gtrsim 10^{10}$ G are somewhat larger than the changes in the spectrum and beaming pattern that occur when going from a purely H model atmosphere to a purely He model atmosphere without any magnetic field. While spectra that assume such strong fields might yield statistically acceptable fits to the *NICER* observations of PSR J0030+0451, it is likely, though not certain, that such fits would, as in the case of the He atmosphere models, yield mass and radius estimates that are so extreme that we would consider these models disfavored.

If the total strengths of the magnetic fields in the atmospheres of the hot spots of PSR J0030+0451 are less than B_c , they are expected to have little effect on the spectrum and shape of its soft X-ray waveform. The agreement of our energy-resolved waveform models with the energy-resolved waveform data obtained using *NICER* are consistent with this interpretation.

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

In modeling the soft X-ray pulse waveform of PSR J0030+0451, we started with the simplest possible description of the heated region(s) on the stellar surface and then, as necessary, considered more complicated descriptions, guided at each step by the differences between the best-fit simpler pulse waveform models and the waveform observed using *NICER*. Along the way, we verified our parameter estimation algorithms by fitting appropriate models to several synthetic waveforms that were generated using the models of the heated regions on the stellar surface that emerged from our fits to the observed waveform, and then comparing the values of the parameters obtained from these fits with the values of these same parameters that were used to generate the synthetic waveforms.

Visual examination of the PSR J0030+0451 waveform observed using *NICER* shows evidence for two peaks, and waveform models constructed using a single hot spot fail to reproduce the qualitative features of the observed waveform. Consequently, we began our modeling of the observed waveform using models that had two uniform circular hot spots with possibly different locations, areas, and temperatures, but these models did not adequately describe the observed waveform. In particular, the (relatively large) spot sizes that were required to reproduce the weak higher harmonic content of the observed waveform appeared to conflict with the (relatively small) spot areas that were required to reproduce the observed X-ray flux.

We therefore considered models with three and four different uniform-temperature circular hot spots, and finally models with two and three different uniform-temperature oval hot spots. In all cases we allowed the spots to overlap in any way that improved the fit; our procedure for treating overlapping spots is described below. This algorithm allowed the heated regions in the model to evolve toward a variety of different complex shapes as the fitting process proceeded, including multiple separate hot spots, multiple-temperature spots in which each spot has hotter and cooler regions, elongated or more complicated configurations produced by two or more partially or completely overlapping or adjacent circular or oval hot spots, crescent-shaped hot regions produced by cold oval spots partially or completely overlapping hot oval spots, and so on. These explorations found no evidence for different temperatures within the same hot spot and no evidence for more than three heated regions on the stellar surface.

We found that a model with two different, non-overlapping, uniform oval spots is preferred over any models with two overlapping spots, adequately describes the observed waveform, and gives a much better fit to the data than a model with two circular spots. The larger east–west extent of the best-fitting oval spots allows them to reproduce the observed weak higher harmonic content of the waveform, while their smaller north–south extent allows them to have the relatively small areas required to reproduce the observed X-ray flux.

To assess whether a model with two different uniform-temperature oval spots is sufficient, we also considered configurations with three different uniform-temperature oval spots. We found that a model with three different, non-overlapping, uniform oval spots is preferred over any models in which the spots overlap, describes the observed waveform adequately, and gives a fit to the data that is slightly, but not significantly, better than the best-fit model with two non-overlapping, uniform oval spots. The two main spots in the best-fit model with three spots are very similar to the spots in the two-spot model, while the third spot has a much higher temperature but a much smaller area than either of the other two spots, is located close to the rotational pole on the far side of the star from the observer, and makes only a very small contribution to the waveform.

The pulse waveforms and the mass and radius estimates given by the best-fit two- and three-spot models are statistically indistinguishable. We therefore concluded that both models adequately describe the observed waveform, that there is no evidence that a more complicated model is required, and that the mass and radius estimates given by these models are reliable. We slightly prefer the radius and mass estimates given by the model with three oval spots, only because the evidence for this model is slightly—although not significantly—greater than the evidence for the model with two oval spots.

In Section 4, we show and discuss the results of several of the analyses of synthetic waveform data that we performed to gain confidence in the accuracy and reliability of our parameter estimation procedure. Specifically, we show the results that we obtained by fitting a model of the waveform produced by two circular spots to synthetic waveform data generated using this model, and by fitting a model of the waveform produced by two oval spots to synthetic waveform data generated using this more complicated model. In Section 5, we describe in detail our systematic approach to modeling the pulse waveform of PSR J0030+0451 and the evidence that led us to focus on the waveform model with three oval spots that we used to estimate the radius and mass of PSR J0030+0451. But first we detail the various waveform models that we constructed, the modeling procedure that led us to consider these models, and the algorithms that we used to estimate the best-fit values of the parameters in these models. We also describe the procedures that we used to evaluate and compare these models, the motivation for our selection of the *NICER* data we analyzed, and our assessment of the effects of the errors in the *NICER* response.

3.3. Waveform Models

In the analyses that we present here, we used pulse waveform models with the primary parameters listed and defined in Table 1. There are 12 primary parameters in the model that has two circular spots (this model assumes $f_1 = 1 = f_2$ and has no parameters for a third spot), 14 primary

Table 1

Primary Parameters of the Pulse Waveform Models Considered in this Work

Parameter	Definition	Assumed Prior
$GM/(c^2R_c)$	Stellar compactness	0.125–0.3125
M	Gravitational mass	1.0–2.4 M_\odot
θ_{c1}	Spot 1 center	0.1–3.1 rad
$\Delta\theta_1$	Spot 1 radius	0–3 rad
f_1	Spot 1 oval ratio	0.1–20
$kT_{\text{eff},1}$	Spot 1 eff. temp.	0.011–0.5 keV
$\Delta\phi_2$	Spot 2 long- itude diff.	0–1 cycles
θ_{c2}	Spot 2 center	0.1–3.1 rad
$\Delta\theta_2$	Spot 2 radius	0–3 rad
f_2	Spot 2 oval ratio	0.1–20
$kT_{\text{eff},2}$	Spot 2 eff. temp.	0.011–0.5 keV
$\Delta\phi_3$	Spot 3 long- itude diff.	0–1 cycles
θ_{c3}	Spot 3 center	0.1–3.1 rad
$\Delta\theta_3$	Spot 3 radius	0–3 rad
f_3	Spot 3 oval ratio	0.1–20
$kT_{\text{eff},3}$	Spot 3 eff. temp.	0.011–0.5 keV
θ_{obs}	Observer inclination	0.1– $\pi/2$ rad
N_{H}	H column density	0– $2.5 \times 10^{20} \text{ cm}^{-2}$
D	Distance	$\exp[-(D - 0.325\text{kpc})^2/2(0.017\text{kpc})^2]$

Note. We assumed the above priors, which are flat in the indicated range except for the distance D , and a rotation frequency of 205.53 Hz as seen by a distant observer (a) when we analyzed the synthetic PSR J0030+0451 waveform generated using two oval hot spots, which had parameter values (see Table 4) chosen so that it would mimic the waveform observed by *NICER*, and (b) when we analyzed the actual PSR J0030+0451 waveform observed by *NICER*. When we analyzed the synthetic waveform that was generated using two circular hot spots, we used the priors listed above for most of the parameters, but different priors for four parameters. This is because the waveform was constructed to mimic the waveform of PSR J0437–4715 observed by *NICER*, and in generating it we used values for these four parameters that are different from those listed in Table 2. In generating this synthetic waveform, we assumed a rotation frequency of 173.6 Hz and flat priors in the intervals listed for the following four parameters: M : 1.0–1.9 M_\odot , θ_{obs} : 0.6–0.9 rad, N_{H} : 0– $2.5 \times 10^{20} \text{ cm}^{-2}$, D : 0.13–0.18 kpc. See the text for further details about the assumed prior probability distributions for each parameter.

parameters in the model that has two oval spots, and 19 primary parameters in the model that has three oval spots.

For the reasons discussed previously, all the waveform models that we present in this report assume that the effective temperature of the emission from a given hot spot does not vary across the spot. The effective temperature listed in Table 1 is the effective temperature that would be measured by a local observer on the stellar surface.

When constructing an oval spot centered at a given colatitude and having a given latitudinal extent, we determined the longitudinal extent it would have if it were instead circular

and then multiplied this extent by the colatitude-independent factor f_1 for spot 1, f_2 for spot 2, and f_3 for spot 3. The possible values of f_1 , f_2 , and f_3 are bounded from above by the requirement that no spot have a longitudinal extent greater than 2π .

We defined the northern hemisphere as the hemisphere containing the sightline to the observer. We defined longitude 0 as the longitude of the center of spot 1; after generating a trial waveform, we rotated its phase to give the best fit to the waveform data (see below).

3.4. Our Waveform Modeling Procedure

As noted earlier, when fitting model waveforms to the observed waveform, we allowed hot spots to overlap one another. This approach allowed the fitting process to create more complicated heated regions, such as multiple isolated spots, spots with hotter cores and cooler annuli, spots with cooler cores and hotter annuli, noncircular heated regions having two different temperatures, and so on, if these configurations were favored by the data. For the reasons described in Section 3.2, we discuss here only our modeling of waveforms with at least two hot spots and at most three.

In the pixels on the stellar surface where spots overlap, our fitting algorithm chooses the temperature of the lowest-numbered spot. For example, if spot 2 and spot 3 overlap, all the pixels in the overlap region are assumed to emit with the effective temperature currently assigned to spot 2. If spot 1, spot 2, and spot 3 all overlap, all pixels in the overlap region are assumed to emit with the effective temperature currently assigned to spot 1.

In addition to the primary parameters listed in Table 1, we introduced two types of ancillary parameters. The first set of ancillary parameters describes the counts in each energy channel that do not come from any of the hot spots. We refer to these as *background* counts. These background counts allow for any unmodulated emission from the star that is not produced by any of the hot spots, any other unmodulated emission from the pulsar system, any emission within the field of view that does not come from the pulsar system, and the instrumental backgrounds. Another ancillary parameter describes the overall starting phase of the pulse waveform. In our fitting procedure, we chose to marginalize over these ancillary parameters, rather than fitting for them explicitly. Separately marginalizing the likelihood over the ancillary parameters is fully justified when, as here, the values of the primary parameters are uncorrelated with the values of the ancillary parameters.

We estimated the number of background counts in each energy channel and marginalized the likelihood with respect to the number of background counts in each channel as follows. We first generated a trial waveform by choosing a trial set of values for all the parameters in the waveform model and a trial starting phase for the waveform. Then, for each energy channel, we independently fit a Gaussian to the curve of the likelihood versus the number of background counts assumed present in that channel. By integrating over each Gaussian, we were able to estimate the number of background counts in each energy channel and marginalize the likelihood over the number of background counts in each energy channel, for each trial waveform. This procedure added one parameter to the waveform model for each energy channel that we considered. As we discuss in detail below, for our final fits to the *NICER* pulse

waveform data we chose to use *NICER* energy channels 40 through 299. This approach therefore added 260 parameters to the waveform model.

We performed a similar procedure to determine the best-fit overall phase of the waveform and marginalize the likelihood over this phase. Namely, for each trial offset of the overall phase (defined as the difference between the longitude of the center of the first spot relative to the longitude of the line of sight to the observer), we computed the likelihood over all energy channels. We then fit a Gaussian to the likelihood as a function of the overall phase, which allowed us to estimate the overall phase. We then marginalized the likelihood with respect to the overall phase by integrating the phase offset over the Gaussian curve, assuming a prior for the overall phase that was uniform from 0 to 1 cycles and zero outside this range. This added one more parameter to the waveform model.

Our procedure for estimating the number of background counts in each *NICER* energy channel assumes that we have no independent knowledge of the number or spectrum of the background counts. In reality, we have some information about astrophysical and sky backgrounds and the backgrounds contributed by solar system, magnetospheric, terrestrial, and International Space Station sources. We have developed codes that can take this information into account, but many of these backgrounds are time-variable and are often not accurately known. The values of the waveform parameters that we derived when we included models of the properties of these backgrounds differ little from the values we found when we used the procedure just described. We have therefore chosen to adopt the conservative approach of using this procedure, which does not rely on any models of the various backgrounds.

3.5. Parameter Estimation and Model Evaluation

Parameter estimation and model comparison. To estimate the values of the parameters in our pulse waveform models and compare the evidence for different models using synthetic or *NICER* data, we used standard Bayesian methods. For a given model \mathcal{M} with parameters θ that have a normalized prior probability density $q(\theta)$, the normalized posterior probability density is given by Bayes' Theorem:

$$P(\mathcal{M}, \theta | \text{data}) = \frac{\mathcal{L}(\text{data} | \mathcal{M}, \theta) q(\theta)}{P(\mathcal{M} | \text{data})}, \quad (1)$$

where $\mathcal{L}(\text{data} | \mathcal{M}, \theta)$ is the likelihood of the data given the model with parameter values θ , and $P(\mathcal{M} | \text{data}) \equiv \int_{\theta} \mathcal{L}(\text{data} | \mathcal{M}, \theta) q(\theta) d\theta$ is the evidence for model \mathcal{M} .

Given a model \mathcal{M} , we can compute the joint posterior probability density distribution of the values of any subset of the parameters θ by marginalizing the other parameters, i.e., by integrating $P(\mathcal{M}, \theta | \text{data})$ over the possible values of all parameters except those in the subset. For example, the 1D probability density distribution of the value of any single parameter can be computed by integrating $P(\mathcal{M}, \theta | \text{data})$ over the possible values of all the other parameters.

The relative probability of two different models is given by their odds ratio, which is the ratio of their evidences (their Bayes factors) multiplied by the ratio of the prior probabilities of the two models.

Sampling methods. We used three statistical sampling methods to estimate the values of the parameters in our various waveform models and compute the evidence for these models:

the publicly available nested sampler MultiNest (MN; Feroz et al. 2009), the parallel-tempering Markov chain Monte Carlo sampler (PT-emcee) that is included in the publicly available emcee package version 2.2.1 (Foreman-Mackey et al. 2013), and a hybrid sampling algorithm (see below) that we call MN+PT-emcee.

The MN and PT-emcee sampling algorithms are complementary: nested samplers are optimized for computing the Bayesian evidence and are therefore well suited for computing the odds ratios of different models, whereas Markov Chain Monte Carlo algorithms are optimized for parameter estimation. We carried out many tests of MN, PT-emcee, and MN+PT-emcee in the context of our waveform-fitting problem, using both synthetic and *NICER* waveform data sets.

For relatively simple waveform models and data sets, such as those with two uniform circular hot spots, we were able to adequately sample the full parameter space using either MN or PT-emcee in a reasonable amount of clock time. For these simpler models and data sets, the posterior probability density distributions given by these two algorithms appeared well converged and were in excellent agreement.

For more complex models and data sets, such as our fits of models with two or three oval spots to synthetic data or to the *NICER* data on PSR J0030+0451, we found that, at least for the values of its sampling parameters that we explored, the MN algorithm abandoned regions of parameter space that contained solutions with log likelihoods comparable to the highest found in previous searches, did not give reproducible results, and—when we used it to analyze synthetic data—produced 1σ and 2σ credible regions that too often did not contain the values of the parameters that had been assumed in generating the synthetic data, even after we ran the code for a very long time.

In contrast, even with poor initial seeds, the PT-emcee algorithm appeared able to sample well the parameter space of complicated models and data sets, gave reproducible results, and produced 1σ and 2σ credible regions that contained the values of the parameters that had been assumed in generating the synthetic data about as often as one would expect, i.e., about 68% of the time for 1σ credible regions and about 95% of the time for 2σ credible regions. We therefore adopted the PT-emcee algorithm for analyzing the more complex models and data sets that we considered. However, with poor initial seeds, PT-emcee took large amounts of clock time to complete each parameter estimation.

In order to speed up parameter estimation using PT-emcee, we sought to choose high likelihood points as initial guesses. Having the MN nested sampling algorithm readily available, we typically used it to provide an initial guess. We tested this hybrid approach, which we refer to as MN+PT-emcee, extensively, and found that it was able to find correct solutions in much less clock time than was required if PT-emcee was not given a high likelihood initial guess. Consequently, we used MN+PT-emcee to estimate the values of the parameters in the models with two and three oval spots, and used the resulting samples to estimate the evidence for these models, given the data.

In our initial MultiNest surveys, we used 1000 active points (essentially the number of parameter combinations being actively updated), an efficiency of 0.01 (the target fraction of new parameter combinations that update the list of active points), and a tolerance of 0.1; for precise definitions of these parameters; see Feroz et al. (2009). Based on our experience

sampling the synthetic and *NICER* waveform data, we limited the number of modes to five for models with two spots and 10 for models with three spots. In our PT-emcee analyses, we used 400 walkers when fitting models with two spots and 1000 walkers when fitting models with three spots, and the default temperature ladder.

Completeness of models. Although our parameter estimations and model comparisons are fully Bayesian, we also performed χ^2 analyses to determine whether our models adequately describe the *NICER* data. In computing χ^2 , we use the expression for model variance originally advocated by Pearson (1900), namely,

$$\chi^2 = \sum_i \frac{(m_i - d_i)^2}{m_i}, \quad (2)$$

where m_i is the number of counts in bin i that are expected in the model, and d_i is the observed number of counts in bin i , rather than the common data variance form in which the m_i in the denominator is replaced by d_i . The model variance form has the advantage that if the counts d_i are drawn from a Poisson distribution with expected values m_i , then the expected value of χ^2 is equal to the number of bins for any values of m_i . This is not true for the data variance form. We note that for our purposes the computation of χ^2 is a one-way test: if the value of χ^2 for a given number of degrees of freedom is large enough to be highly improbable then we are warned that our model of the spots and/or the instrument is likely to be incomplete, but if the value of χ^2 is reasonably probable we cannot conclude that the model is correct.

3.6. Treatment of Instrumental Uncertainties and Bias

The method that we used to make the final parameter estimates we report here takes into account the substantially larger uncertainties in the *NICER* response below 0.4 keV (channel 40) and the evidence we found that including data from these channels might bias our results on PSR J0030+0451.

As we discuss in Section 5.2, we found that including data from channels below channel 40 biases waveform-based estimates of the distance to PSR J0030+0451 toward values much larger than the independently known distance. The source of this bias is currently unknown. Since its origin is unknown, it could also bias waveform-based estimates of other parameters, including the radius and mass of the star. Given this risk, we judged it important to address this issue.

It might be possible to gain a better understanding of the origin of this distance bias if we also had reliable, independent estimates of the values of other stellar and waveform parameters, but for PSR J0030+0451, we only have a reliable independent estimate of the distance. This provides too little information to identify the source(s) of this bias or to devise sophisticated corrections. We therefore considered only simple ways to address this bias. We also wanted to be careful not to do anything that might bias other, even more important model parameters.

As was mentioned in Section 2.3, the *NICER* response below channel 40 is more uncertain than the response above channel 40. It was therefore natural to explore the effect of ignoring the data in channels below channel 40. As we show in detail in Section 5.2, discarding the data in these low channels eliminates the bias in our waveform-based estimates of the

distance to PSR J0030+0451. It is our hope that discarding these data also reduces any potential biases in our estimates of the other parameters in the waveform model. Because this data selection procedure is so simple and does not touch the data that we do use, we think it is less likely to bias other waveform parameters than a more complex procedure.

In analyzing the *NICER* data on PSR J0030+0451, we found no detectable modulation in the energy channels above channel 299, and we therefore did not include data from these channels in our analyses of the *NICER* data on PSR J0030+0451.

One would expect that discarding the data in channels 25–39 would reduce the precisions of the parameter estimates. We investigated this possibility by comparing the precisions of the parameter estimates obtained by analyzing the data in channels 40–299 produced by a synthetic waveform that is similar to the PSR J0030+0451 waveform observed by *NICER* with the precisions obtained by analyzing the data in channels 25–299 produced by the same synthetic waveforms and found that discarding the data in channels 25–39 did not significantly reduce the precisions of parameter estimates (see Section 4.2). We also compared the posterior probability distributions for the radius and mass of PSR J0030+0451 obtained using the *NICER* data in channels 40–299 with the probability distributions obtained using the data in channels 25–299 and found no evidence that the precisions of the parameter estimates were reduced by discarding the data in channels 25–39 (see Section 5.2).

A remaining question is whether the fits of our waveform models to the *NICER* waveform data on PSR J0030+0451 are affected by the unavoidable systematic errors present in our model of the *NICER* response in the energy channels 40–299 that we used. As we described in Section 2.3, the *NICER* response in these channels was calibrated by comparing the results of *NICER* observations of the Crab Nebula, which is traditionally assumed to have a power-law spectrum, with the power-law spectrum of the nebula determined by observations made using other instruments. When *NICER* observations of the Crab Nebula spectrum are compared with the spectrum obtained by folding the power law determined by other instruments through the nominal *NICER* response, the residuals in the energy channels we used in this study (channels 40–299) are $\lesssim 5\%$ (Ludlam et al. 2018). As mentioned in Section 2.3, these residuals are likely due to imperfect modeling of the microphysics of the concentrator optics and astrophysical features, such as oxygen edges and emission lines.

We explored the likely effect of deviations of the nominal *NICER* response from the true *NICER* response curve by multiplying the nominal effective area curve of *NICER* by the ratio between the count rate as a function of energy given by the assumed spectrum of the Crab Nebula and the count rate as a function of energy measured by *NICER* (Ludlam et al. 2018). We then used this modified effective area curve to analyze both synthetic waveform data constructed to mimic the *NICER* data on PSR J0030+0451, and the actual *NICER* waveform data on PSR J0030+0451.

We found that modifying the nominal effective area in this way had a negligible effect on the quality of our fits to synthetic waveform data and to the actual waveform data on PSR J0030+0451. For example, the probability of the χ^2 obtained by comparing the energy-resolved waveform given by the best-fit model with two oval spots to the *NICER* data in energy

Table 2
Fits to Synthetic Data Generated Using Two Circular Spots

Parameter	Assumed Value	Median	-1σ	$+1\sigma$	-2σ	$+2\sigma$
R_e (km)	13.0	13.435	12.248	14.722	11.191	15.929
$GM/(c^2 R_e)$	0.1636	0.163	0.149	0.174	0.137	0.182
M (M_\odot)	1.44	1.472	1.289	1.667	1.151	1.849
θ_{c1} (rad)	0.6283	0.531	0.456	0.608	0.400	0.680
$\Delta\theta_1$ (rad)	0.01	0.011	0.010	0.012	0.009	0.013
$kT_{\text{eff},1}$ (keV)	0.231	0.229	0.223	0.233	0.219	0.237
θ_{c2} (rad)	2.077	2.096	1.959	2.260	1.871	2.401
$\Delta\theta_2$ (rad)	0.33	0.358	0.322	0.397	0.288	0.446
$kT_{\text{eff},2}$ (keV)	0.0578	0.056	0.054	0.059	0.052	0.061
$\Delta\phi_2$ (cycles)	0.5625	0.562	0.556	0.567	0.550	0.571
θ_{obs} (rad)	0.733	0.774	0.662	0.894	0.583	0.996
N_H (10^{20} cm^{-2})	2.0	2.266	2.132	2.398	1.999	2.536
D (kpc)	0.156	0.165	0.146	0.176	0.133	0.179

Note. Results obtained using our MN+PT-emcee sampling algorithm to fit a waveform model with two possibly different and overlapping uniform circular spots to the data in *NICER* energy channels 25–299 from a synthetic waveform generated assuming two different, non-overlapping uniform circular spots.

channels assigned to 32 phase bins changed from 0.105 to 0.112 when we used the modified effective area curve instead of the nominal effective area curve. This indicates that the current systematic errors in the nominal effective area of *NICER* in channels 40–299 have only a minor effect on our fits.

Given these results, we chose to ignore all events in the channels below channel 40 and use the current nominal *NICER* effective area curve for channels 40–299, both when analyzing the actual *NICER* data on PSR J0030+0451 and when analyzing synthetic *NICER* data.

4. Analysis of Synthetic Pulse Waveform Data

In order to verify that our parameter estimation procedures give correct posterior probability distributions for the range of pulse waveform models that we used to analyze the *NICER* waveform data on PSR J0030+0451, we applied these same procedures to synthetic waveform data sets that we created to mimic the *NICER* data on PSR J0030+0451, and checked that our procedures give appropriate parameter estimates and credible regions.

As we explained in the previous section, we chose to discard the data on PSR J0030+0451 in the *NICER* energy channels below channel 40 and above channel 299. In order to verify that discarding the data in the channels below channel 40 would not bias our estimates of the waveform parameters, and to assess the effect of discarding these data on the precisions of our estimates, we first analyzed the data in channels 25–299 given by a particular synthetic waveform and then analyzed only the data in channels 40–299 given by the same synthetic waveform (we did not consider any channels below channel 25, because channel 25 was the lowest channel included in the cleaned *NICER* data set). In each case, we analyzed the synthetic data using the MultiNest sampling algorithm and also using either the PT-emcee sampling algorithm or our hybrid MN+PT-emcee algorithm.

By (1) fitting models having two possibly different and overlapping circular spots to synthetic data generated assuming two different, non-overlapping circular spots and (2) fitting models having two possibly different and overlapping oval spots to synthetic data generated assuming two different, non-overlapping oval spots, we were able to verify that our parameter estimation procedures yield appropriate posterior probability density distributions for these synthetic waveforms. Specifically, for these data sets the MultiNest and MN+PT-emcee sampling algorithms both produced posterior distributions for the waveform parameters that are fully consistent with the values of the parameters that were assumed when the synthetic waveforms were generated. We view the agreement of the results given by these two different sampling algorithms with each other and with the parameters assumed in generating the synthetic data as evidence that both algorithms give correct results when used to analyze synthetic waveforms like these, when the sampling is complete.

By comparing the results we obtained by analyzing the synthetic waveform data in *NICER* energy channels 40–299 with the results that we obtained by analyzing the data in channels 25–299, we were also able to assess the effect on the precisions of our parameter estimates of using only the data in energy channels 40–299. We found that using this reduced data set had a negligible effect on the sizes of the credible regions.

We now provide the details of these tests.

4.1. Analysis of Synthetic Waveforms Generated Using Two Uniform Circular Spots

The first analysis of synthetic data we present here shows the results we obtained by fitting a 12-parameter waveform model with two possibly different and overlapping uniform circular spots to a synthetic waveform generated assuming two different, non-overlapping uniform circular spots, using our MN+PT-emcee sampling algorithm. Table 2 lists the values of

Table 3

Verification of Fits to Synthetic Data Generated by Two Circular Spots

Energy Channels	$\pm 1\sigma$ (68.3%)	$\pm 2\sigma$ (95.4%)	$\pm 3\sigma$ (99.7%)
25–299	10	12	12
40–299	7	11	12

Note. The number of values of the waveform parameters assumed in generating the synthetic waveform that are within the indicated 1D credible intervals derived from our fit of the 12-parameter waveform model with two possibly different and overlapping circular spots to the synthetic waveform generated assuming two different, non-overlapping circular spots, using the data in the indicated energy channels.

the waveform parameters that we assumed when we generated the synthetic waveform data. We note that the temperature we assumed for one of the two spots when we generated this synthetic waveform is substantially less than the temperatures of either of the two main spots we found when we analyzed the *NICER* data on PSR J0030+0451. We show our results for this synthetic waveform because it was the synthetic waveform that was selected—in advance of the analyses of the PSR J0030+0451 waveform—to verify the parameter estimation procedures that would be used by the *NICER* team to analyze pulse waveforms. The results obtained by others on the team when they analyzed this waveform are consistent with the results we obtained when we analyzed it (Bogdanov et al. 2019b).

Table 2 shows the results that we obtained by fitting the waveform model with two possibly different and overlapping uniform circular spots to synthetic waveform data generated assuming two different, non-overlapping uniform circular spots. The table lists the median value of each model parameter computed using its 1D posterior probability distribution and the boundaries of the $\pm 1\sigma$ and $\pm 2\sigma$ credible intervals for each parameter.

We found that when this waveform model was fit to the synthetic waveform data in *NICER* energy channels 25–299, the values of 10 of the 12 parameters assumed in generating the synthetic waveform are within the relevant $\pm 1\sigma$ credible intervals for these parameters while the values of the remaining two parameters (θ_{c1} and N_H) assumed in generating the synthetic waveform are within their $\pm 2\sigma$ credible intervals. We found that when this model was instead fit only to the synthetic data in *NICER* energy channels 40–299, the assumed values of seven of the 12 parameters were within the relevant $\pm 1\sigma$ credible intervals, the assumed values of four others were within the relevant $\pm 2\sigma$ credible intervals, and the assumed value of the other parameter was within the relevant $\pm 3\sigma$ credible interval. Thus, for both energy channel choices, the 1D posterior probability distributions derived for the waveform parameters are consistent with the values of the parameters assumed when the synthetic waveform was generated. These results are summarized in Table 3.

Figure 1 shows that the joint probability density distribution for M and R_e that was obtained by fitting this model to the synthetic waveform data in *NICER* energy channels 25–299 is consistent with the values of the stellar mass and radius assumed when the synthetic waveform was generated.

The panels in the top row of Figure 2 show the 1D probability distributions of M , R_e , and D obtained by fitting this model to the synthetic waveform data in *NICER* energy




Figure 1. 68% (inner) and 95% (outer) contours of the joint probability density distribution of M and R_e obtained by fitting a model waveform with two possibly different uniform circular hot spots to a synthetic waveform that assumed two different uniform circular hot spots. This analysis used the synthetic waveform data in *NICER* energy channels 25–299. The joint probability density of M and R_e obtained from this fit is consistent with the M and R_e values assumed when the synthetic data were generated, which are indicated by the star.

channels 25–299, whereas the panels in the bottom row show the results obtained by fitting to the data in channels 40–299. This figure illustrates two noteworthy aspects of these distributions.

First, these 1D probability distributions are consistent with each other and with the values of M , R_e , and D that were assumed in generating the synthetic waveform. The widths of the distributions for M and R_e obtained using only the data in channels 40–299 underestimate the true uncertainties in the estimates of these parameters, because the temperature assumed for the lower-temperature spot when generating this synthetic waveform was sufficiently low that discarding the data in energy channels 25–39 narrowed these distributions. Our tests (see, e.g., Figure 4 below) have shown that—in contrast to the reduced widths of the M , R_e , and D distributions obtained from this synthetic waveform when the data in channels 25–39 are discarded—discarding the data in these channels has no significant effect on the distributions of M , R_e , and D obtained by analyzing the *NICER* data on PSR J0030+0451, which has higher-temperature hot spots.

Second, although the 1D probability distributions for D obtained by fitting this model to data selected using both these energy-channel criteria are consistent with the value of D assumed in generating the synthetic waveform data (the cumulative probability near the assumed value of D is substantial for both channel selections and fits with wider priors show that the probability density eventually decreases with increasing distance), these results show that pulse waveform analysis does not determine D very precisely. These analyses demonstrate the importance of using the precise and accurate independent measurement of D that is available for

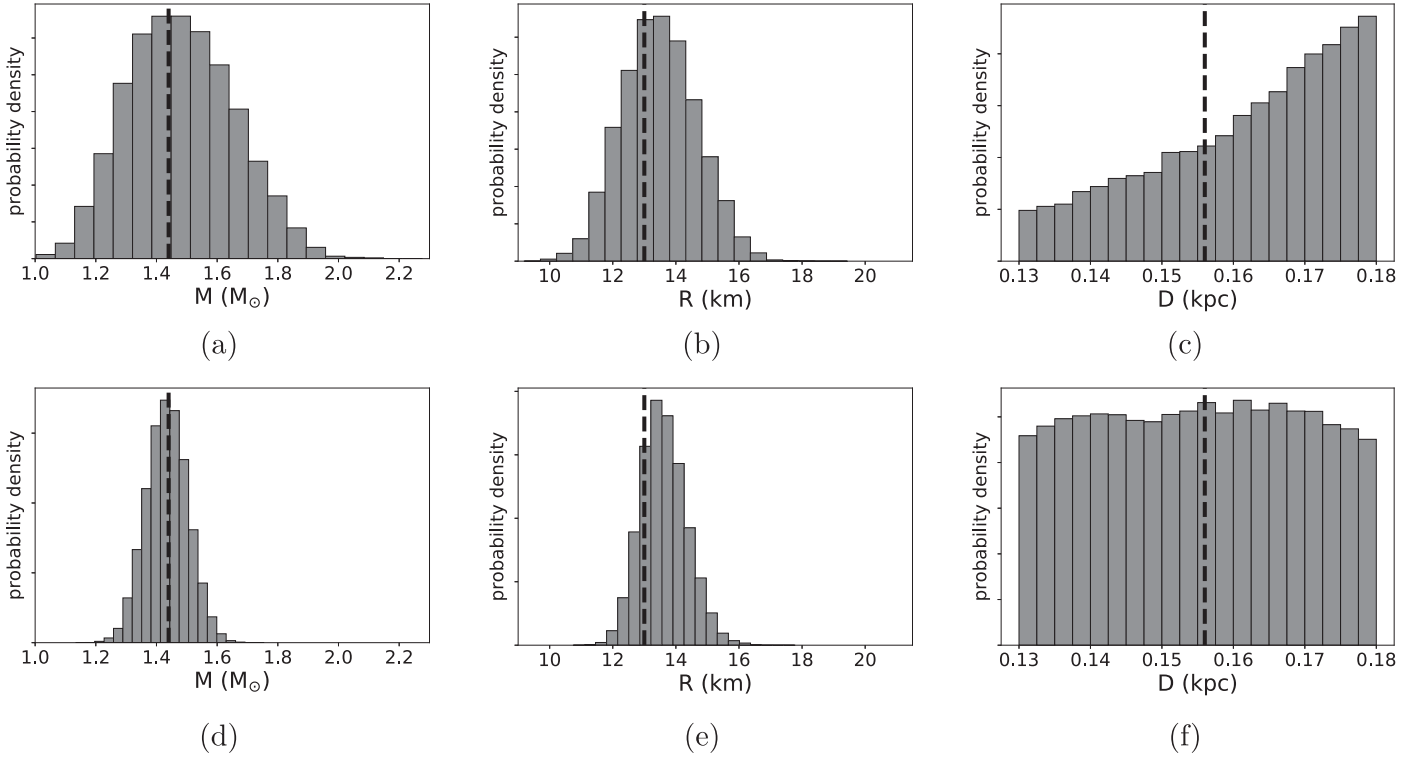


Figure 2. Posterior probability density distributions for M , R_e , and D obtained by fitting the waveform model with two possibly different uniform circular hot spots to synthetic waveform data generated using a model that assumed two different uniform circular hot spots. Top row: results obtained by fitting the data in *NICER* energy channels 25–299. Bottom row: results obtained by fitting the data in *NICER* energy channels 40–299. The vertical dashed line in each panel shows the value of that waveform parameter that was assumed when generating the synthetic waveform data.

Table 4
Fits to Synthetic Data with Two Oval Spots

Parameter	Assumed Value	Median	-1σ	$+1\sigma$	-2σ	$+2\sigma$
R_e (km)	13.49	15.997	14.497	17.471	13.044	18.798
$GM/(c^2 R_e)$	0.1503	0.145	0.136	0.153	0.129	0.161
M (M_\odot)	1.374	1.565	1.386	1.750	1.225	1.918
θ_{e1} (rad)	2.251	2.247	2.181	2.319	2.113	2.393
$\Delta\theta_1$ (rad)	0.026	0.024	0.021	0.028	0.018	0.033
f_1	9.14	7.356	5.589	9.558	4.603	11.742
$kT_{\text{eff},1}$ (keV)	0.1151	0.114	0.111	0.117	0.109	0.119
θ_{e2} (rad)	2.442	2.459	2.401	2.516	2.336	2.567
$\Delta\theta_2$ (rad)	0.031	0.028	0.025	0.032	0.022	0.038
f_2	16.23	17.465	14.672	19.288	11.941	19.907
$kT_{\text{eff},2}$ (keV)	0.1164	0.115	0.113	0.117	0.111	0.119
$\Delta\phi_2$ (cycles)	0.457	0.455	0.452	0.457	0.450	0.459
θ_{obs} (rad)	0.939	0.931	0.848	1.003	0.766	1.071
N_{H} (10^{20} cm^{-2})	0.084	0.163	0.051	0.355	0.007	0.654
D (kpc)	0.352	0.350	0.315	0.369	0.284	0.374

Note. Results obtained using our MN+PT-emcee sampling algorithm to fit a waveform model with two possibly different and overlapping uniform oval spots to *NICER* energy channels 40–299 of a synthetic waveform generated assuming two different, non-overlapping uniform oval spots.

Table 5

Verification of Fits to Synthetic Data Generated by Two Oval Spots

Energy Channels	$\pm 1\sigma$ (68.3%)	$\pm 2\sigma$ (95.4%)	$\pm 3\sigma$ (99.7%)
25–299	10	13	14
40–299	12	14	14

Note. The number of values of the waveform parameters assumed in generating the synthetic waveform that are within the indicated 1D credible intervals derived from our fit of the 14-parameter waveform model with two possibly different and overlapping oval spots to the synthetic waveform generated assuming two different, non-overlapping oval spots, using the data in the indicated energy channels.

PSR J0030+0451 (see below) when analyzing the pulse waveform of this pulsar.

4.2. Analysis of Synthetic Waveforms Generated Using Two Uniform Oval Spots

The second analysis of synthetic data we present here shows the results we obtained by fitting a 14-parameter waveform model with two possibly different and overlapping uniform oval spots to a synthetic waveform generated assuming two different, non-overlapping uniform oval spots, using our MN+PT-emcee sampling algorithm. Table 4 lists the values of the waveform parameters that we assumed when we generated the synthetic waveform data.

The results that we obtained by fitting this waveform model to the synthetic waveform data in *NICER* energy channels 25–299 and 40–299 are shown in Table 4, which lists the median value of each model parameter computed using its 1D posterior probability distribution and the boundaries of the $\pm 1\sigma$ and $\pm 2\sigma$ credible intervals for each parameter.

We found that when the model waveform was fit to the synthetic waveform data in *NICER* energy channels 25–299, the values of 10 of the 14 parameters assumed in generating the synthetic waveform are within the $\pm 1\sigma$ credible intervals for these parameters, the values assumed for three others are within the relevant $\pm 2\sigma$ credible intervals, and the value assumed for the remaining parameter is within the relevant $\pm 3\sigma$ credible interval. When this model waveform was fit only to the synthetic data in *NICER* energy channels 40–299, the assumed values of 12 of the 14 parameters fall within their $\pm 1\sigma$ credible intervals and the assumed values of the other two parameters fall within their $\pm 2\sigma$ credible intervals. The waveform parameter values that were assumed when generating the synthetic waveform data are all consistent with the 1D posterior probability densities obtained by analyzing the data using either energy cut. These results are summarized in Table 5.

Figure 3 shows that the joint probability density distribution for M and R_e obtained by fitting this model to the synthetic waveform data in *NICER* energy channels 25–299, and by fitting it to the synthetic data in channels 40–299, are both consistent with the values of the stellar mass and radius assumed when the synthetic data was generated. Figure 3 also demonstrates that discarding the data in *NICER* energy channels 25–39 does not significantly degrade the precision of the M and R_e estimates.

The panels in the top row of Figure 4 show the 1D probability distributions of M , R_e , and D obtained by fitting this model to the synthetic waveform using the data in *NICER*




Figure 3. Plot of the 68% (inner) and 95% (outer) contours of the joint probability density distribution of M and R_e obtained by fitting a model waveform with two possibly different and overlapping uniform oval hot spots to a synthetic waveform that assumed two different, non-overlapping uniform oval spots. The solid black lines show the contours obtained when the synthetic waveform data in *NICER* energy channels 25–299 were used, whereas the dashed blue lines show the contours obtained when the synthetic waveform data in energy channels 40–299 were used. Both joint probability distributions are consistent with the M and R_e values assumed when the synthetic data were generated, which are indicated by the star. These results show that using only the data in energy channels 40–299 neither biases the M and R_e estimates nor significantly reduces their precisions.

energy channels 25–299, whereas the panels in the bottom row show the results obtained using only the data in energy channels 40–299. We note several aspects of these distributions.

First, these 1D probability distributions are consistent with each other and with the values of M , R_e , and D that were assumed in generating the synthetic waveform. Second, these distributions are not significantly affected by using only the data in energy channels 40 and above. Finally, although the 1D probability distributions for D obtained by fitting this model to these data are consistent with the value of D assumed in generating the synthetic waveform (the cumulative probability near the assumed value of D is substantial for both energy cuts and fits with wider priors show that the probability density eventually decreases with increasing distance), waveform analysis again does not determine D very precisely. This further emphasizes the value of using the precise and accurate independent measurement of D (see below) when analyzing the *NICER* waveform data on PSR J0030+0451.

5. Analysis of the *NICER* Pulse Waveform Data

In this section we describe our estimation of the radius and mass of PSR J0030+0451 using *NICER* observations of its soft X-ray pulse waveform. We first describe our modeling of this waveform and detail the basis for our decision to use the data in *NICER* energy channels 40–299. We then present our radius and mass estimates, and the evidence for the validity of the model that we used to make these estimates. Finally, we

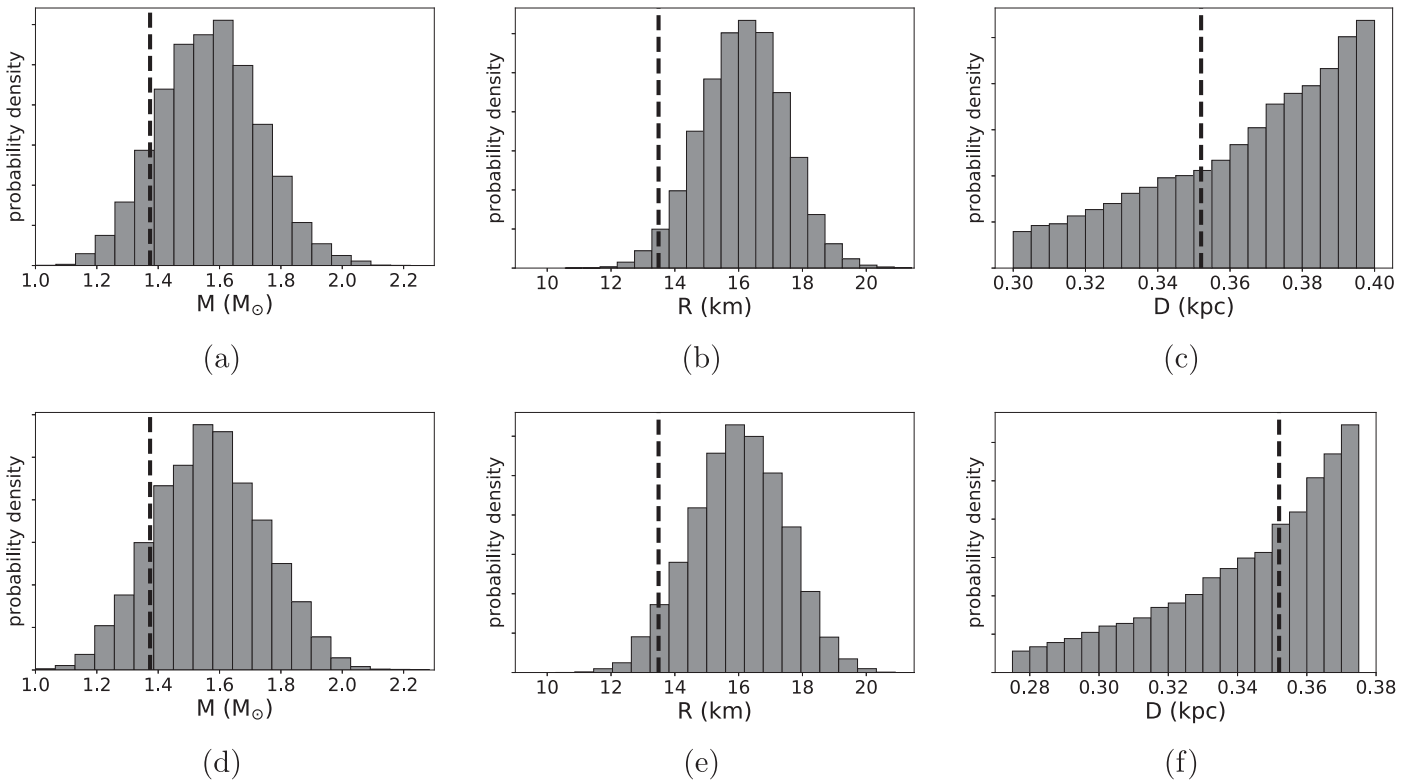


Figure 4. Posterior probability density distributions for the M , R , and D estimates obtained by fitting the waveform model with two possibly different and overlapping uniform oval hot spots to synthetic waveform data generated using a model that assumed two different, non-overlapping uniform oval hot spots. Top row: results obtained by fitting the data in *NICER* energy channels 25–299. Bottom row: results obtained by fitting the data in *NICER* energy channels 40–299. The vertical dashed line in each panel shows the value of that waveform parameter that was assumed when generating the synthetic waveform data.

discuss our choices of waveform-phase resolution, the sizes and locations of the hot spots given by our modeling, the best-fit colatitude of the observer, and the combined, waveform- and phase-integrated X-ray spectrum of the hot spots given by our analysis.

As we have previously described in Section 3.5, we explored a variety of algorithms for sampling the parameter space of waveform models but found the best results using the hybrid sampling algorithm that we refer to as MN+PT-emcee. Therefore, in this section we present only results obtained using this algorithm.

5.1. Modeling the Waveform of PSR J0030+0451

Our algorithm for developing a successful model of the heated regions on the surface of PSR J0030+0451 was to proceed systematically from the simplest possible model to more complicated models, guided at each step by the differences between the waveform given by the best-fit simpler models and the waveform observed using *NICER*, and ending when adding additional hot spots or complexity fails to increase significantly the evidence for the model. The sequence of models that we considered in our quantitative modeling is presented in Table 6. We now discuss this sequence.

As mentioned in Section 3.2, visual examination of the PSR J0030+0451 waveform shows evidence for two peaks. As a result, waveform models constructed using a single hot spot fail to reproduce the qualitative features of the observed waveform. We therefore began our quantitative modeling of the observed waveform using models with two uniform circular hot spots, allowing the spots to have different locations, areas, and

temperatures. We also allowed the two circular spots to overlap partially or completely. The sampling algorithm therefore explored configurations in which the two circular spots (1) are disjoint, with the same or different temperatures; (2) partially overlap, potentially producing a somewhat oval-shaped uniform-temperature spot if the two component spots have the same temperature, or a spot with a complicated boundary and a complex temperature structure, if the two component spots have different temperatures; or (3) fully overlap, producing (3a) a single, uniform-temperature circular spot, if one of the spots completely covers the other, (3b) a spot with a centered or off-center core-annulus structure in which the core is hotter or colder than the annulus, if the two spots have different temperatures, or (3c) a circular spot partially surrounded by a hotter or colder crescent-shaped region. None of these models produced waveforms that adequately describe the observed waveform. Among other problems, the (relatively large) spot sizes required to reproduce the weak higher harmonic content of the observed waveform appeared to conflict with the (relatively small) spot areas required to reproduce the observed X-ray flux. We therefore considered more complicated heated regions.

In the second step of our exploration of waveform models, we considered models with three and four uniform circular hot spots. In all cases we allowed the spots to overlap in any way that improved the fit. This algorithm allowed the heated regions in the model to evolve toward a variety of different complex shapes as the fitting process proceeded (see Section 3.2). These explorations found no significant evidence for configurations with more than three circular hot spots on the stellar surface, for configurations in which three hot spots were near one another,

Table 6
Summary of Pulse Waveform Models Considered and the Results

Waveform Model	Results
Two circular spots	Strongly disfavored The required areas of the spots appear to conflict with the harmonic structure of the waveform
Three or four circular spots	Disfavored No evidence for overlapping spots No evidence for configurations with more than three spots No evidence for more than one temperature in any of the heated regions The required areas of the spots appear to conflict with the harmonic structure of the waveform
Two oval spots	Strongly favored over models with two circular spots Two different non-overlapping oval spots adequately describe the observed waveform data The best-fit temperatures of the two spots are almost the same One spot is slightly elongated in the east–west direction The other spot is highly elongated in the east–west direction No evidence that a more complicated configuration is needed to describe the observed waveform
Three oval spots	Strongly favored over models with two circular spots Three different non-overlapping oval spots adequately describe the observed waveform data Radius and mass estimates are statistically indistinguishable from the estimates using the model with two oval hot spots Two of the spots are almost identical to the two spots in the best-fit model with two oval spots The third spot makes almost no contribution to the waveform Slightly but not significantly higher evidence than the model with two oval spots No evidence that a more complicated configuration is needed to describe the observed waveform

Note. See the text for details about the assumed prior probability distributions for each parameter.

or for overlapping or adjacent hot spots with significantly different temperatures, such as one would expect if there were spots on the stellar surface that had regions with different temperatures or a significant temperature gradient. We again found that the larger spot sizes required to reproduce the weak higher harmonic content of the waveform appeared to conflict with the smaller sizes required to reproduce the observed X-ray flux. This conflict led us to consider uniform-temperature oval hot spots as the next step in our sequence of trial models.

The first oval hot-spot model that we considered examined heated regions that could be modeled using two uniform-temperature oval spots. We again allowed the two spots to have different properties and to overlap one another, partially or completely. The sampling algorithm could have collapsed this model to a model with two uniform-temperature circular spots, if that was preferred by the data, but it was not. The sampling algorithm therefore explored configurations in which the two oval spots (1) are disjoint, with the same or different temperatures; (2) partially overlap, potentially producing a uniform-temperature spot with a complicated boundary, if both spots have the same temperature, or a spot with a complicated boundary and a complex temperature structure, if the two spots have different temperatures; or (3) fully overlap, producing (3a) a single, uniform-temperature oval spot, if one of the spots completely covers the other, (3b) a possibly oval spot with a potentially oval centered or off-center core-annulus structure in which the core is

hotter or colder than the annulus, if the two spots have different temperatures, or (3c) a possibly oval spot partially surrounded by a hotter or colder crescent-shaped region.

Our analysis showed that a configuration with two non-overlapping, uniform oval spots with almost the same temperatures is preferred over any models with two overlapping spots, adequately describes the observed waveform, and is strongly favored over the best-fit model with two circular spots. Indeed, the log evidence for the best-fit model with two oval spots was 11.6 larger than the log evidence for the best-fit model with two circular spots. The fit with two uniform oval spots favors a configuration in which both spots are elongated in the east–west direction and both have almost the same temperature. The larger east–west extent of the best-fitting oval spots allows them to successfully reproduce the observed weak higher harmonic content of the waveform while the smaller north–south extent allows them to have the relatively small areas required to reproduce the observed X-ray flux.

To assess whether a model with two uniform-temperature oval spots is sufficient, we also considered configurations with three, possibly different, uniform-temperature oval spots. We found that a configuration with three different non-overlapping oval spots also describes the observed waveform adequately. This model is again strongly favored over the best-fit model with two uniform circular spots: the log evidence for the best-fit model with three oval spots was 13.3 greater than the log

evidence for the best-fit model with two circular spots. We again found no evidence for more than one temperature in any of the heated regions. Two of the three spots in the best-fitting three-spot model are almost identical to the two spots in the best-fitting two-spot model. Like the two spots in the best-fitting two-spot model, the two larger spots in the best-fitting three-spot model are elongated in the east–west direction and have almost the same temperatures. The third spot in the best-fitting three-spot model has a much higher temperature but a much smaller area than either of the other two spots, is located close to the rotational pole on the far side of the star from the observer, and makes only a very small contribution to the waveform. The waveform produced by the best-fitting model with three oval spots is very similar to the waveform produced by the model with two oval spots.

The best-fitting models with two and three oval spots both provide acceptable fits to the phase-channel and bolometric waveform data in the χ^2 sense, and are not statistically distinguishable from each other. Most importantly, the best-fit values of the stellar radius and mass given by the model with three oval spots are almost identical to, and are statistically indistinguishable from, the best-fit values of the stellar radius and mass given by the model with two oval spots. The evidence for the best-fit model with three oval spots is slightly, but not significantly, greater than the evidence for the model with two oval spots.

Based on these results, we conclude that the models with two and three oval spots both provide adequate descriptions of the observed waveform, that there is no evidence that a more complicated model is required to describe the observed waveform, and that the mass and radius estimates inferred using these two models are reliable. We slightly prefer, and therefore report here, the radius and mass estimates given by the model with three oval spots, only because the evidence for this model is slightly—although not significantly—greater than the evidence for the model with two oval spots.

These considerations and results are summarized in Table 6. We now provide the details.

5.2. Choice of Energy Channels

The lowest-energy *NICER* data that was included in the standard pipeline data set that we analyzed was in energy channel 25 (see Section 2.2). As noted previously (see Section 3.6), we found no detectable modulation of the X-ray flux in *NICER* energy channels 300 and above. We therefore considered in our analyses only the data in channels 25–299.

The calibration errors in the data from energy channels 40–299 are relatively small. For example, the residuals between the Crab Nebula spectrum determined from *NICER* measurements and the Crab Nebula spectrum determined from measurements made using other instruments are $\lesssim 5\%$ for channels 40–299 (see Sections 2.3 and 3.6; see also Ludlam et al. 2018). In contrast, there are still substantial and poorly understood errors in the calibration of energy channels 25–39. We were therefore concerned that including data from channels 25–39 in our analysis would bias our estimates of the mass and radius of PSR J0030+0451.

The distance to PSR J0030+0451 has been accurately measured to be 0.325 ± 0.009 kpc, independent of any waveform analyses (see Arzoumanian et al. 2018a). We therefore decided to investigate the effect of including or not including the data in channels 25–39 on estimates of the distance made using

the PSR J0030+0451 waveform. We found that if we used the *NICER* waveform data in channels 25–299, thereby including the data in channels 25–39, and assumed a wide, flat prior for the distance, all our analyses gave posterior probability distributions for the distance that peaked at distances substantially larger than the independently measured distance of 0.325 kpc and excluded 0.325 kpc with high confidence. This is illustrated by panel (c) in Figure 5, which shows the 1D posterior distributions for M , R_e , and D obtained by fitting the waveform model with two possibly different and overlapping uniform oval hot spots to the *NICER* data in channels 25–299. In contrast, when we used only the well-calibrated data in channels 40–299, we found posterior probability distributions for the distance that naturally peak close to the independently measured distance. This is illustrated by panel (f) in Figure 5, which shows the 1D posterior distributions for M , R_e , and D obtained by fitting the waveform model with two possibly different and overlapping uniform oval hot spots to the *NICER* data in channels 40–299.

These results suggested to us that including the poorly calibrated data in channels 25–39 in the analysis was likely to be biasing the distance estimate and could also be biasing the estimates of M and R_e , about which we have no independent information. We therefore thought it was important to address the effects of the apparent calibration errors in channels 25–39, but we wanted to do this in a way that would not itself bias our results. Consequently, we decided that, rather than attempting to correct for the calibration errors in channels 25–39, which are poorly understood at present, the safest approach would be to discard the data in all these channels. This would prevent the calibration errors in these channels from biasing our M and R_e estimates, but it could in principle increase the uncertainties in our estimates of these parameters because the amount of data we used would be reduced.

Comparison of the results obtained by analyzing synthetic waveform data in *NICER* energy channels 25–299 and 40–299 that were generated assuming two different, non-overlapping uniform circular spots (see Section 4.1) shows that using only the data in the reduced energy range does not significantly bias the estimates of M and R_e or decrease the precisions of the estimates (compare panels (a) and (d) and panels (b) and (e) of Figure 2)). (We note that for this synthetic waveform, the temperature assumed for one of the hot spots was substantially lower than any of the temperatures inferred for the hot spots of PSR J0030+0451. Consequently, the cut at channel 40 discarded much more of the spectrum of this waveform than it does for the spectrum of PSR J0030+0451 observed by *NICER*.)

Similarly, comparison of the results obtained by analyzing synthetic waveform data in *NICER* energy channels 25–299 and 40–299 that were generated assuming two different, non-overlapping uniform oval spots (see Section 4.2) shows that using only the data in the reduced energy range does not significantly bias the estimates of M and R_e or decrease the precisions of the estimates (compare panels (a) and (d) and panels (b) and (e) of Figure 4).

Finally, comparison of the results obtained by analyzing the *NICER* data on the waveform of PSR J0030+0451 in energy channels 25–299 and 40–299 shows that using only the data in the reduced energy range does not significantly decrease the precisions of the estimates of M and R_e (compare panels (a) and (d) and panels (b) and (e) of Figure 5).

Taking into account all these results, we decided to use only the data in energy channels 40–299 to estimate the values of the

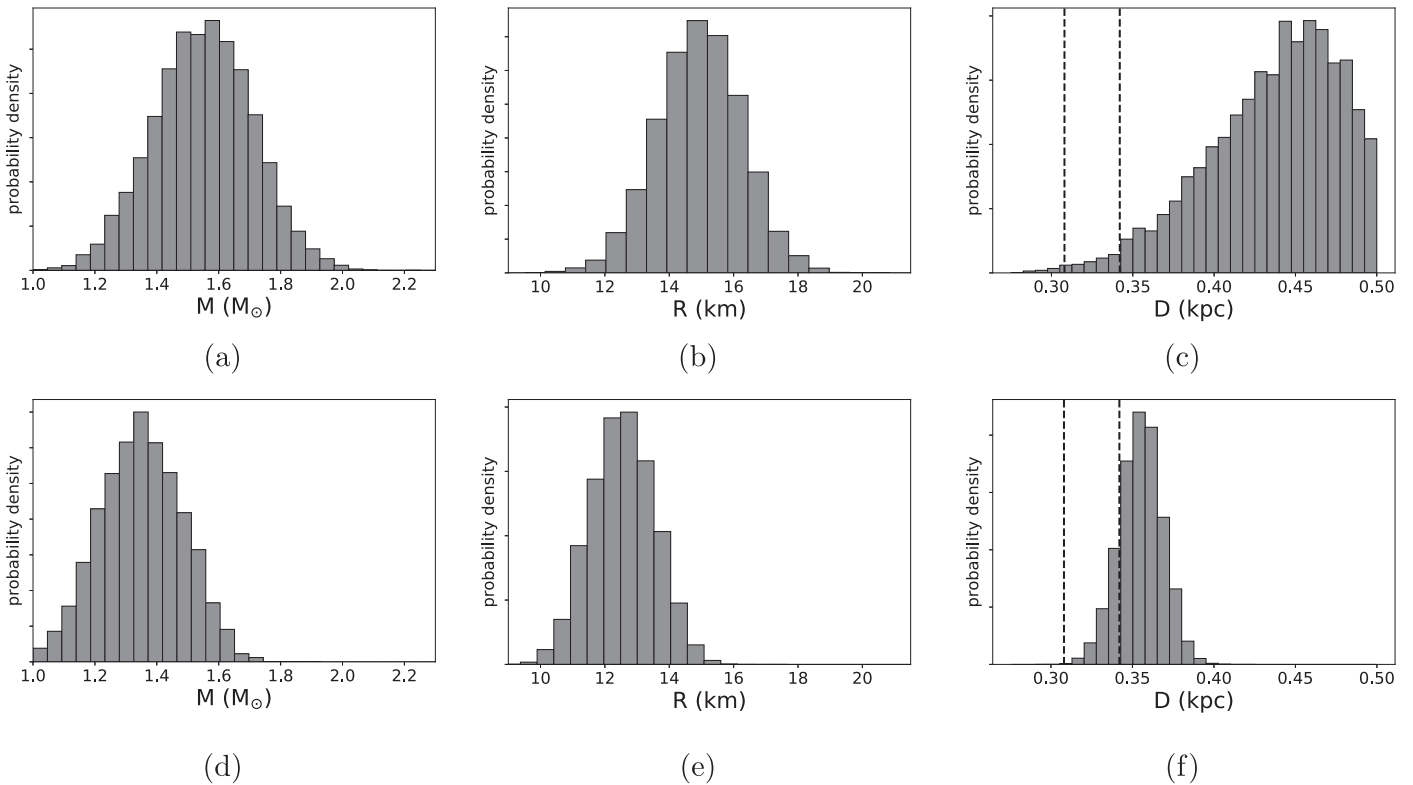


Figure 5. 1D posterior probability distributions for M , R_e , and D estimates obtained by fitting the waveform model with two possibly different uniform oval hot spots to two different sets of *NICER* data on PSR J0030+0451. Both fits assumed a distance prior flat between 0.275 and 0.5 kpc and zero outside this interval. The independently measured distance to PSR J0030+0451 is 0.325 ± 0.009 kpc (see the text). The vertical dashed lines in panels (c) and (f) indicate the uncertainty in D obtained by combining the uncertainty in the measurement of D and the uncertainty in D that corresponds to the uncertainty in the effective area of *NICER* (see the text for details). Top row: posteriors obtained using the data in *NICER* energy channels 25–299. Note that the distance posterior, which is plotted in panel (c), peaks at ~ 0.45 kpc and excludes the independently measured distance with high confidence. Bottom row: posteriors obtained using the data in *NICER* energy channels 40–299. In contrast to the distance posterior obtained using the data in *NICER* energy channels 25–299, the distance posterior obtained using this data set, which is plotted in panel (f), peaks at a distance that is close to the independently measured distance.

parameters in our models of the PSR J0030+0451 pulse waveform. Consequently, in the remainder of this Letter we present only results obtained using the data in these energy channels. Given that the independently measured distance of 0.325 ± 0.009 kpc is very precise, we imposed a Gaussian prior on the distance that uses this measurement and an uncertainty that takes into account the uncertainty in this measurement and the global (in energy) uncertainty in the effective area of the *NICER* instrument (see below).

5.3. The Mass and Radius of PSR J0030+0451

In this section, we present and discuss the results of our analysis of the *NICER* data on PSR J0030+0451, focusing on our estimates of its mass and radius. For the reasons described in Section 5.2, we present only the results we obtained by fitting the *NICER* data in energy channels 40–299. When estimating the values of all the parameters except the distance, we assumed prior probability distributions constant within the intervals listed in Table 1 and zero outside these intervals. When estimating the distance, we assumed a Gaussian prior. We chose the width of the Gaussian by combining linearly, rather than quadratically, the 0.009 kpc uncertainty in the measured distance to PSR J0030+0451 (see Arzoumanian et al. 2018a) and the estimated uncertainty in the overall normalization of the *NICER* effective area, converted to an uncertainty in the distance. We assumed a 5% uncertainty in the overall normalization of the *NICER* effective area (see

Section 3.6). This translates into a distance uncertainty of 2.5%. The distance to PSR J0030+0451 has been independently measured to be 0.325 kpc (Arzoumanian et al. 2018a). Consequently, a 1σ uncertainty of 2.5% in the distance is ≈ 0.008 kpc, yielding a total 1σ width for the Gaussian of 0.009 kpc + 0.008 kpc = 0.017 kpc.

We considered two waveform models, one with two possibly different and overlapping uniform-temperature oval spots, and one with three possibly different and overlapping uniform-temperature oval spots. The model with two oval spots has 14 primary parameters. These are listed in Table 7, along with their median values and the $\pm 1\sigma$ and $\pm 2\sigma$ boundaries of their 1D credible regions, computed using the 1D posterior probability density distribution we found for each parameter. The model with three oval spots has 19 primary parameters. These are listed in Table 8, along with their median values and the $\pm 1\sigma$ and $\pm 2\sigma$ boundaries of their 1D credible regions, computed using the 1D posterior probability density distribution we found for each parameter.²⁰ Both fits favored models in which the spots do not overlap.

Table 9 compares the median values of R_e and M and the boundaries of the $\pm 1\sigma$ credible regions computed from the posterior probability density distributions obtained by fitting the two models described here.

²⁰ Samples from the full posterior probability distribution obtained by fitting the model with three uniform-temperature oval spots to the waveform data are available at <https://zenodo.org/record/3473466>.

Table 7
Fits to *NICER* Data with Two Oval Spots

Parameter	Median	-1σ	$+1\sigma$	-2σ	$+2\sigma$	Best Fit
R_e (km)	13.271	12.115	14.578	11.042	15.968	13.643
$GM/(c^2R_e)$	0.160	0.152	0.169	0.143	0.176	0.161
M (M_\odot)	1.442	1.282	1.619	1.141	1.802	1.488
θ_{c1} (rad)	2.251	2.165	2.334	2.084	2.419	2.232
$\Delta\theta_1$ (rad)	0.035	0.030	0.041	0.026	0.047	0.031
f_1	5.347	3.950	6.981	3.136	8.362	6.024
$kT_{\text{eff},1}$ (keV)	0.117	0.114	0.120	0.111	0.122	0.119
θ_{c2} (rad)	2.417	2.333	2.495	2.245	2.560	2.394
$\Delta\theta_2$ (rad)	0.033	0.028	0.039	0.025	0.044	0.029
f_2	15.490	12.317	18.396	10.020	19.748	17.744
$kT_{\text{eff},2}$ (keV)	0.117	0.115	0.119	0.112	0.122	0.119
$\Delta\phi_2$ (cycles)	0.459	0.457	0.461	0.455	0.463	0.458
θ_{obs} (rad)	0.848	0.751	0.951	0.660	1.043	0.827
N_{H} (10^{20} cm^{-2})	0.266	0.072	0.593	0.011	1.08	0.047
D (kpc)	0.327	0.309	0.345	0.293	0.361	0.332

Note. 1D credible regions, and best fit, obtained by fitting the model with two possibly different uniform oval spots to channels 40–299 of the *NICER* data on PSR J0030+0451.

Figure 6 compares the 1D posterior probability density distributions for M , R_e , and D . Note that the waveform data modifies the posterior probability distributions for the distance only slightly compared to the Gaussian prior that we assumed based on the independently measured distance to PSR J0030+0451 and the uncertainty in the *NICER* effective area (see Section 5.2), showing that the distance estimates obtained by fitting both models to the waveform data are consistent with the measured distance.

Figure 7 compares the joint posterior probability density distributions for R_e and M that were obtained by fitting the two models to the *NICER* waveform data.

Figure 8 compares the bolometric waveforms given by the best-fit waveform models with two and three oval spots. The waveform components are very similar for the two models, and the full waveforms are almost identical.

All these comparisons show that the models with two and three uniform-temperature oval spots give mutually consistent credible regions for all the parameters that they have in common, most notably for the mass and radius of the pulsar. The log evidence for the best fit of the model with three oval spots is higher than the log evidence for the best fit of the model with two oval spots, but only by 1.7, which is less than the estimated uncertainty in the log evidence. Thus, both models are comparably good. This consistency suggests that, at least for the types of waveform models considered here, no further complexity is needed to adequately describe the observed waveform. The 68% credible intervals for M and R_e given by the fit of the model with three oval spots to the *NICER* data from PSR J0030+0451 are 1.30–1.59 M_\odot and 11.96–14.26 km.

Our results for mass and radius of PSR J0030+0451 may be compared with the recent results of Steiner et al. (2018), who examined the evidence from neutron stars that are transient

X-ray sources and conclude that the radius of a 1.4 M_\odot neutron star is most likely between 10 and 14 km, which is consistent with our estimates of the mass and radius of PSR J0030+0451.

Our results may also be compared with those of Nättilä et al. (2017), who estimated that the mass of the (neutron star) X-ray burst source 4U 1702–429 is between 1.6 and 2.2 M_\odot and that its circumferential radius is between 12.0 and 12.8 km (with 68% credibility). Although this estimate of the mass of 4U 1702–429 is higher than our estimate of the mass of PSR J0030+0451, for most neutron star EoS currently under consideration, the radius of a neutron star depends only weakly on its mass in this mass range, in which case this estimate of the radius of 4U 1702–429 is consistent with our estimate of the radius of PSR J0030+0451.

The full posterior probability density distributions for the parameters in the fit of the model with two oval spots to the *NICER* waveform data are shown in Figure 15 in Appendix A. The full posterior probability density distributions for the parameters in the fit of the model with three oval spots are shown in Figure 16 in Appendix B. Among other things, these “corner” plots show that there is a correlation between estimates of the stellar radius and the distance, but that this correlation is fairly weak. Thus, any errors in the data that bias the estimate of the distance may not significantly bias the estimate of the radius. There are indications that our sampling of the posterior distribution of the stretching parameter for the second oval spot did not explore a wide enough range, but this probably did not significantly affect our results, because there is no correlation between the value of this parameter and the values of the mass, radius, or distance parameters.

5.4. Adequacy of the Models

We performed χ^2 analyses (see Equation (2)) to determine whether the models we have used to estimate the radius and mass

Table 8
Fits to *NICER* Data with Three Oval Spots

Parameter	Median	-1σ	$+1\sigma$	-2σ	$+2\sigma$	Best Fit
R_c (km)	13.019	11.959	14.255	10.938	15.500	13.466
$GM/(c^2R_c)$	0.163	0.154	0.171	0.144	0.179	0.156
M (M_\odot)	1.443	1.299	1.594	1.164	1.745	1.423
θ_{c1} (rad)	2.270	2.179	2.357	2.093	2.442	2.330
$\Delta\theta_1$ (rad)	0.036	0.031	0.040	0.028	0.045	0.032
f_1	5.352	4.364	6.502	3.568	7.664	5.335
$kT_{\text{eff},1}$ (keV)	0.117	0.114	0.120	0.110	0.122	0.113
θ_{c2} (rad)	2.417	2.341	2.486	2.252	2.540	2.446
$\Delta\theta_2$ (rad)	0.033	0.029	0.038	0.025	0.043	0.029
f_2	15.769	13.017	18.498	10.923	19.79	16.588
$kT_{\text{eff},2}$ (keV)	0.115	0.112	0.118	0.107	0.121	0.105
$\Delta\phi_2$ (cycles)	0.460	0.458	0.463	0.456	0.466	0.463
θ_{c3} (rad)	2.988	2.865	3.063	2.691	3.094	3.056
$\Delta\theta_3$ (rad)	0.056	0.021	0.103	0.004	0.168	0.087
f_3	1.215	0.493	2.096	0.161	3.511	1.253
$kT_{\text{eff},3}$ (keV)	0.239	0.143	0.354	0.029	0.470	0.209
$\Delta\phi_3$ (cycles)	0.420	0.378	0.534	0.071	0.932	0.427
θ_{obs} (rad)	0.878	0.769	0.973	0.675	1.051	1.012
N_{H} (10^{20} cm^{-2})	0.244	0.082	0.441	0.011	0.723	0.187
D (kpc)	0.327	0.307	0.347	0.290	0.365	0.317

Note. 1D credible regions, and best fit, obtained by fitting the model with three possibly different uniform oval spots to channels 40–299 of the *NICER* data on PSR J0030+0451.

of PSR J0030+0451 adequately describe the *NICER* data. Specifically, we compared our best-fit energy-resolved model waveforms with the energy-resolved waveform data (i.e., the pulse-phase–energy–channel data) collected by *NICER*, and compared the energy-integrated (bolometric) waveforms given by our best-fit energy-resolved waveform models with the bolometric waveforms observed by *NICER*, using the values of χ^2 given by these comparisons. (In principle, we could also compare the pulse-phase-integrated spectra given by our models with the pulse-phase-integrated spectrum observed by *NICER* by computing the relevant values of χ^2 , but our procedure for modeling the observed unmodulated counts—see Section 3.4—guarantees a nearly perfect description of the photon energy spectrum, so this comparison would be uninformative.)

Assigning the *NICER* data to 32 phase bins and 260 energy channels, the best-fit energy-resolved waveform model with two oval spots gives a χ^2 of 8204.68 when compared with this data set, which consists of the number of counts in each of $32 \times 260 = 8320$ phase-energy bins. As discussed in Section 3.3, the model waveform with two oval spots has 14 primary parameters, 260 ancillary parameters related to the non-star emission, and one additional parameter that describes its overall phase, yielding a total of $8320-260-14-1 = 8045$ degrees of freedom. The resulting $\chi^2/\text{degrees of freedom}$ (dof) is therefore $8204.68/8045$. If this model is correct, the probability of finding a value of χ^2/dof this

large or larger is 0.104. Thus, according to the χ^2 test this model provides an acceptable description of this data set.

Again assigning the *NICER* data to 32 phase bins and 260 energy channels, the best-fit energy-resolved waveform model with three oval spots gives a χ^2 of 8188.99 when compared with the *NICER* data binned in this way. This model has 19 primary parameters, yielding a total of 8040 degrees of freedom. The resulting χ^2/dof is therefore $8188.99/8040$. If this model is correct, the probability of finding a χ^2/dof this large or larger by chance is 0.120, so the fit of this model to this data set is also acceptable, according to the χ^2 test. In Figure 9, we show the value of χ in each of the 8320 phase-energy bins, for this fit. No patterns in the values of χ are evident as a function of phase or energy, which is what one would expect for a good fit.

When the bolometric waveforms predicted by the models with two oval spots and three oval spots that best fit the *NICER* phase-channel data with 32 phases are compared with the 32-phase bolometric waveform constructed using the *NICER* data, the values of χ^2/dof are considerably larger. When the bolometric waveform given by the model with two oval spots is compared with the 32-phase bolometric waveform constructed using the *NICER* data, the value of χ^2 is 40.6. There are $32-14-1-1 = 16$ dof in this 32-phase bolometric data. The χ^2/dof is therefore $40.6/16$, which has a probability 6.3×10^{-4} of occurring by chance. This probability is low enough that it indicates that this

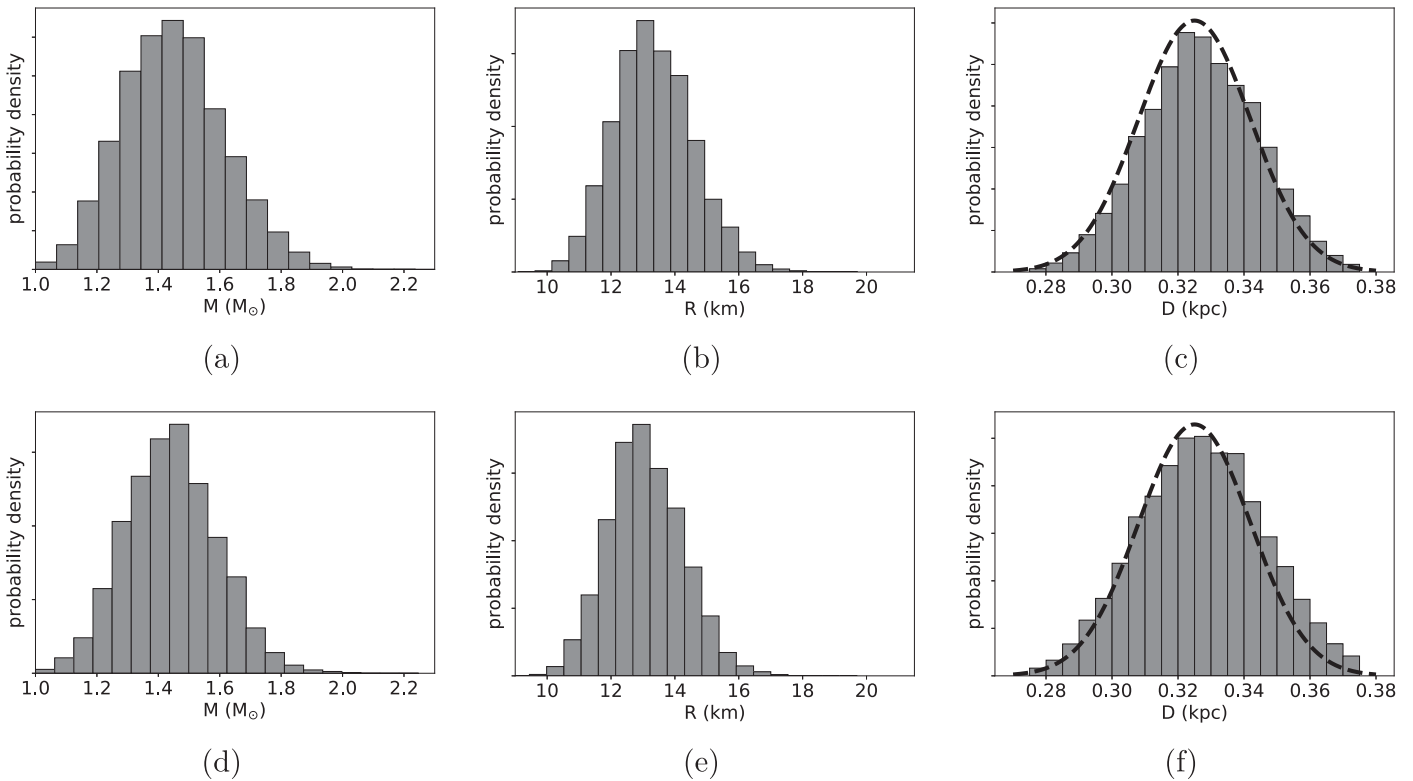


Figure 6. Comparison of the 1D posterior probability density distributions of the stellar mass (panels (a) and (d)), stellar radius (panels (b) and (e)), and the distance to the pulsar (panels (c) and (f)) given by the best fits of the waveform model with two (top row of panels) and three (bottom row of panels) oval spots. Both models were fit to the *NICER* data in energy channels 40–299, assigned to 32 phase bins. The dashed lines in panels (c) and (f) show the Gaussian prior that was used for the distance estimate (see the text). The agreement of the distributions given by the two models is excellent.

Table 9
Comparison between Two-oval and Three-oval Fits to *NICER* Data

Fit	$-1\sigma, R$	Median R_c	$+1\sigma, R$	$-1\sigma, M$	Median M	$+1\sigma, M$
Two ovals	12.118	13.274	14.582	1.282	1.442	1.619
Three ovals	11.962	13.022	14.259	1.299	1.443	1.594

Note. Comparison of the radius (in km) and mass (in solar masses) estimates given by the fits of the models with two and three oval spots to the *NICER* data on PSR J0030+0451.

model provides a description of this data set that is at least incomplete. When the bolometric waveform given by the model with three oval spots is compared with the 32-phase bolometric waveform constructed using the *NICER* data, the value of χ^2 is 41.7. There are $32-19-1-1 = 11$ dof in this 32-phase bolometric data. The χ^2/dof is therefore $41.7/11$, which has a probability 2×10^{-5} of occurring by chance. This probability again indicates that this model provides a description of this data set that is at least incomplete.

Noting these low probabilities, we performed exploratory fits specifically designed to minimize the value of χ^2 obtained when the 32-phase bolometric waveforms predicted by these two energy-resolved waveform models are compared with the 32-phase bolometric waveform constructed from the *NICER* data. The probabilities of the models given by these fits were not significantly larger than the values given by the best fits of these models to the full phase-channel data.

Next, we constructed 64-phase bolometric waveforms using the *NICER* data and compared these with the 64-phase bolometric waveforms predicted by the best fit of the 32-phase

energy-resolved waveform model with two oval spots to the 32-phase energy-resolved *NICER* waveform data, re-fitting the phase shift and the number of phase-independent background counts in each energy channel. The resulting minimum value of χ^2/dof is $58.64/48$, which has a chance probability of 0.140, indicating that this model provides an acceptable description of the 64-phase bolometric waveform data. When the predictions of this model for the 64-phase energy-resolved waveform are compared with 64-phase energy-resolved *NICER* data, the resulting χ^2/dof is $16363.9/16365$, which has a probability of 0.501, much higher than the probability we found when we divided the data into only 32 phase bins, and indicating that this energy-resolved waveform model also provides an acceptable description of the energy-resolved *NICER* waveform data with 64 phase bins.

Finally, we compared the 64-phase bolometric waveforms constructed using the *NICER* data with the 64-phase bolometric waveforms predicted by the energy-resolved waveform model with three oval spots that best fits the 32-phase energy-resolved *NICER* waveform data (again re-fitting the phase shift and the number of phase-independent counts in each energy channel).




Figure 7. Comparison of the joint posterior probability density distributions for M and R_e given by the best fits of the waveform model with two (panel (a)) and three (panel (b)) uniform-temperature oval spots. The inner contour shown in each panel contains 68.3% of the posterior probability, whereas the outer contour contains 95.4%. The color indicates the credibility in standard deviations of each point in the posterior probability density distribution. Again, the agreement of the distributions given by the two models is excellent.




Figure 8. Comparison of the bolometric waveforms given by the best-fit waveform models with two (panel (a)) and three (panel (b)) oval spots. The solid curves show the full waveforms; the dashed curves show the contributions to the full waveform made by the individual hot spots. The components that generate the full waveforms are very similar for the two models.

The results are shown in Figure 10. The minimum value of the bolometric χ^2/dof is 59.6/43, which has a probability of 0.0473, indicating that this 64-phase model also provides an acceptable description of the 64-phase bolometric waveform data. When the predictions of this best-fit model for the energy-resolved waveform with 64 phase bins are compared with the 64-bin energy-resolved *NICER* waveform data, the resulting χ^2/dof is 16347.5/16360, which has a probability of 0.526, again much higher than the probability that we found when we divided the data into only 32 phase bins, and indicating that this

energy-resolved waveform model provides an acceptable description of the energy-resolved *NICER* waveform data with 64 phase bins.

These results indicate that our best-fit models with two and three oval spots provide good descriptions of the *NICER* waveform data at high phase resolutions, and that the radius and mass estimates inferred from them are therefore credible. Why the bolometric waveforms given by the models that best fit the 32- and 64-phase energy-resolved *NICER* waveform data differ from the 32-phase bolometric waveforms constructed




Figure 9. Value of χ in each of the 8320 phase-energy bins (32 phase bins and 260 energy bins), for the best fit of our energy-resolved waveform model with three oval spots to the *NICER* data grouped in 32 phase bins. The energy channel numbers are plotted on a logarithmic scale. No patterns in the values of χ are evident as a function of phase or energy, which is what one would expect for a good fit.




Figure 10. Top: comparison of the 64-phase bolometric waveform constructed using the *NICER* data on PSR J0030+0451 with the 64-phase bolometric waveform model given by the 32-phase energy-resolved waveform model with three oval spots that best fits the 32-phase energy-resolved waveform data. The dashed blue line shows the fitted unmodulated background that was added to the counts produced by the three hot spots as part of the fitting procedure (see Section 3.4). The zero of time of the model bolometric waveform was adjusted to minimize the value of χ^2 ; this adjustment was +0.49 cycles. Bottom: the resulting value of χ as a function of phase. The χ^2/dof is 59.6/43, which has a probability of 0.0473.

from the *NICER* data but agree with 64-phase bolometric waveforms is unclear, but is most likely due to temporal fluctuations in the *NICER* data that are not yet understood. This question deserves further study.

5.5. Other Aspects of the Models

5.5.1. Locations of the Hot Spots

Figure 11 compares the best-fit locations and shapes of the hot spots obtained by fitting the models with two and three oval spots to the *NICER* waveform data. Both hot spots in the model with two oval spots and all three hot spots in the model with three oval spots are located in the southern rotational hemisphere, well away from the sightline to the observer (recall our convention that

the sightline to the observer defines the northern hemisphere). Fits that force any of the spots to be in the same hemisphere as the sightline to the observer are strongly disfavored. As we noted earlier, when we discussed the phase-channel χ^2 values for the fits of both models to the *NICER* waveform data, these spot locations provide formally excellent fits.

The fundamental reason spot locations in the southern hemisphere are favored appears to be that the modulation fraction of the waveform observed by *NICER* is high and its harmonic content is substantial. Both waveform properties favor spot locations in the southern hemisphere, because when the spots are visible for only a small fraction of a rotational cycle, both the modulation amplitude of the waveform and the strengths of the higher harmonics of the spin frequency are larger. Spots located in the northern hemisphere are visible for most of each rotational cycle and therefore produce waveforms with smaller modulation amplitudes and weaker high harmonics. In contrast, radio observations as well as joint fitting of radio and γ -ray profiles suggest geometries in which one of the two radio pulses comes from the northern hemisphere (Johnson et al. 2014). However, these studies assume a star-centered dipolar field, whereas the *NICER* results suggest that the magnetic field of PSR J0030+0451 is not a centered dipole (see also Lockhart et al. 2019). Although the locations of the soft, thermal X-ray emission and the nonthermal radio emission need not be the same, this difference should be investigated further.

5.5.2. Colatitude of the Observer

One recent analysis of γ -ray observations of PSR J0030+0451 suggests that the colatitude θ_{obs} of the observer lies between 0.942 and 1.222 rad (the boundaries of the $\pm 1\sigma$ interval; see Chang & Zhang 2019). The lower end of this interval is consistent with the $\pm 1\sigma$ range of the probability density distributions for θ_{obs} that we infer from the fit of our model with two oval spots and the fit of our model with three oval spots, which is encouraging.

5.5.3. Spectrum of the Emission from the Hot Spots

An independent estimate of the total X-ray emission from PSR J0030+0451 can be made using observations made with

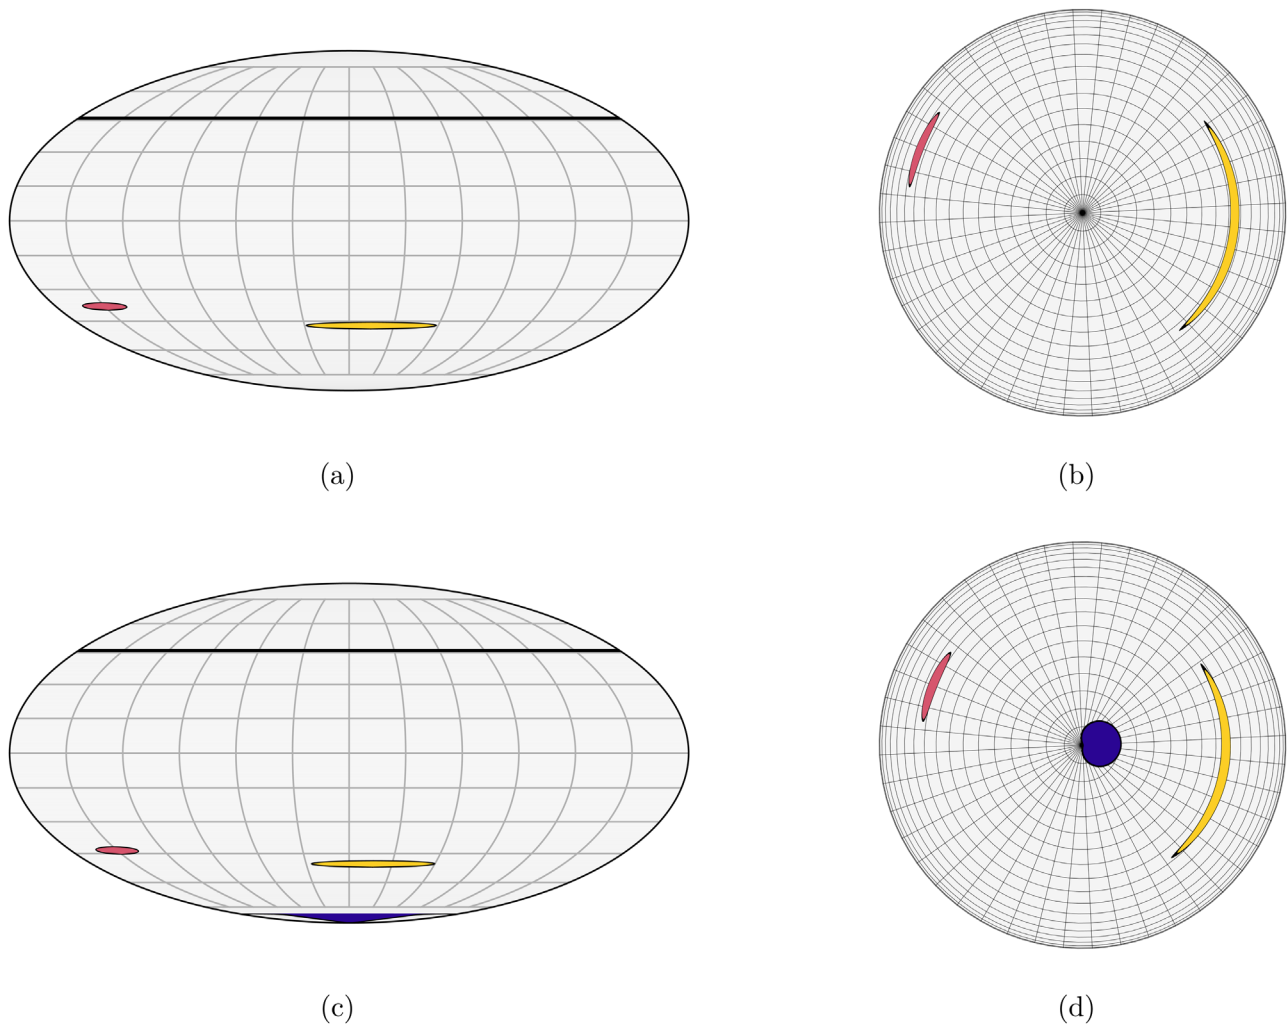


Figure 11. Locations, shapes, and sizes of the hot spots in the best-fit waveform models with two oval spots (panels (a) and (b)) (see Table 7) and three oval spots (panels (c) and (d)) (see Table 8). Panels (a) and (c) show equal-area projections, centered on the rotational equator. Panels (b) and (d) are views from the south pole. The cooler main spot is indicated by yellow, the hotter main spot is indicated by red, and in the three-spot model, the hottest spot is indicated by blue. For both fits, the horizontal line shows the inferred colatitude of the observer. Clearly, both spots in the model with two oval spots and the two main spots in the model with three oval spots are very similar in location, size, and shape; the third spot in the three-spot model has a very small area and makes only a minor contribution to the waveform.

XMM-Newton (Bogdanov & Grindlay 2009). Although the *XMM-Newton* EPIC MOS1/2 observations have substantially fewer source counts, they have a much lower background in the point-source spatial extraction region compared to the non-imaging *NICER* data, which means that to first order all of the counts detected by *XMM-Newton* come from the star rather than from unassociated sources. If the phase-integrated data from *XMM-Newton* are fit using a two-temperature non-magnetic hydrogen atmosphere model (as in Bogdanov & Grindlay 2009) and the predictions of the model are folded through the *NICER* response matrix, then we obtain an estimate for the total number of *NICER* counts from the star and the spectral shape that we expect to see in *NICER*. This predicted spectral flux can be compared with the total spectral flux expected from all the hot spots in our best-fit models of the pulse waveform emission. Figure 12 shows this comparison for our models with two oval spots and three oval spots. In both cases, the combined spectral flux expected from the hot spots at low photon energies falls short of the spectral flux predicted by the *XMM-Newton* observations.

A first thought would be that the missing emission might be unpulsed thermal emission from a substantial fraction of the

stellar surface. There are, however, several serious difficulties with such an interpretation. First, the emission would have to be almost exactly axisymmetric around the stellar rotation axis in order to avoid generating detectable pulsed emission. This forces one to consider emission patterns that are highly tuned: filled circular emitting regions around one or both rotation poles, annular emitting regions centered around one or both rotation axes, or some combination of these would be required to avoid producing detectable flux modulation. Second, if the emission is thermal, the total area of the axisymmetric emitting region(s) would have to be a very small fraction of the stellar surface.

We illustrate these difficulties by an example. In order to make up the observed deficit in the flux at low energies and not produce detectable modulation, a circular spot centered on the north rotational pole (the pole nearer the observer), would have to have an angular radius of just 0.075 rad, assuming thermal emission at the best-fit effective temperature of $kT_{\text{eff}} \approx 0.075$ keV. Such a spot would also have to be very nearly circular and centered on the north rotational pole: a deviation of more than ~ 0.01 rad would cause a flux modulation that would be inconsistent with the observed waveform. Because of its smaller projected area, a circular spot centered on the south rotational pole would have to




Figure 12. Comparison of the predicted total count spectrum that should be observed from PSR J0030+0451 by *NICER* (solid black line), based on the *XMM-Newton* observations of PSR J0030+0451, with the total count spectrum expected from all the hot spots in the best-fit model with two (dashed blue line) and three (solid red line) oval hot spots. The dashed black lines indicate the $\pm 1\sigma$ uncertainty in the count spectrum predicted from the *XMM-Newton* observations. They are a linear combination of the statistical uncertainty in the *XMM-Newton* observations and an estimate of the systematic error in the *NICER* calibration. The total emission predicted in the lower-energy channels using the emission observed by *XMM-Newton* is greater than the total emission expected from all the hot spots in the best-fit hot spot models. This indicates that the emission provided by the hot spots in these models does not account for all of the emission.

have an angular radius ~ 0.7 rad, about 10 times larger than a spot at the north rotational pole, but would have to be even more precisely circular and centered than a spot at the north rotational pole: a deviation of more than ~ 0.005 rad would cause a flux modulation inconsistent with the observed waveform.

Other thermal emission geometries with the same projected area are theoretically possible (e.g., one or more very thin axisymmetric annuli), but these seem to us even more fine-tuned. The spectrum could also be produced by processes in which the peak in the spectrum required to make up the deficit does not reflect the temperature of the emitting plasma, but is instead created by a steep decrease of the optical depth of the emitting region with increasing photon energy. High-harmonic cyclotron emission in a region where the optical depth to cyclotron absorption decreases with increasing photon energy is one such possibility (Psaltis & Lamb 1999), but usually requires a population of higher-energy electrons, in which case the optically thick (Rayleigh–Jeans) portion of the spectrum would have to have a very small area, in order to avoid producing too much emission.

It does appear worthwhile to explore other possibilities. For example, the spot pattern is clearly not that of a centered dipole. If this means that the total magnetic field at the surface of the star has significant multipolar components, and is therefore much stronger than the nominal dipole component of $\sim 2 \times 10^8$ G, then perhaps cyclotron emission could contribute to the missing emission.

6. Implications for the Dense Matter EoS

The joint posterior probability distribution we have obtained for the mass and radius of PSR J0030+0451 provides




Figure 13. Comparison of the mass–radius relations for a wide variety of EoS computed exactly for a stellar spin rate of 200 Hz (solid curves) and approximately using the TOV equation (dotted curves). The EoS are (in order of increasing radius at $1.5 M_\odot$) HPLS1 (Hebeler et al. 2013), BBB2 (the higher mass model that uses the Paris two-body interaction potential from Baldo et al. 1997), APR (Model A18+ $\delta(v)$ +UIX* from Akmal et al. 1998), and HPLS2 (Hebeler et al. 2013). For the $\sim 1.4 M_\odot$ estimated mass of PSR J0030+0451, the exact radii for a spin rate of 200 Hz differ by less than 1% from the approximate radii given by the TOV equation. This difference is far smaller than the uncertainty in the estimate of the radius of PSR J0030+0451 we have obtained from analyzing its soft X-ray waveform.

additional constraints on the EoS of cold catalyzed matter, through the effect of the EoS on the stellar structure. The EoS can be expressed as a relation between the pressure p and the total mass-energy density ρ , or a similar relation between p and a different thermodynamic variable.

Here we use the Tolman–Oppenheimer–Volkoff (TOV) structure equation (Tolman 1939; Oppenheimer & Volkoff 1939)

$$\frac{dp}{dr} = -\frac{(\rho c^2 + p)(M(<r) + 4\pi r^3 p/c^2)}{r(c^2 r/G - 2M(<r))}, \quad (3)$$

where r is the circumferential radius and $M(<r)$ is the gravitational mass inside r . The TOV equation assumes that the star is cold enough that its EoS can be treated as barotropic, which is expected to be an excellent approximation for PSR J0030+0451. The TOV equation also assumes that the star is nonrotating and spherically symmetric. Figure 13 shows that this is an excellent approximation for PSR J0030+0451, which has a rotation frequency of 205 Hz. For a given central mass-energy density, the compactness ratio $GM/(R_e c^2)$ is even less sensitive to the rotation rate than either the mass or the radius considered separately.

Given $p(\rho)$ and a central density, the TOV equation can be integrated to yield the gravitational mass M and circumferential radius R_e of a star.

There are numerous papers that apply various Bayesian techniques to the problem of constraining a parameterization of $p(\rho)$ using various astronomical measurements (e.g., Agathos et al. 2015; Lackey & Wade 2015; Alvarez-Castillo et al. 2016; Riley et al. 2018; Greif et al. 2019). Here we use the approach of Miller et al. (2019), who emphasize using full posterior




Figure 14. Constraints on the EoS (pressure vs. equivalent number density) of cold, catalyzed matter with and without taking into account the *NICER* measurements of the mass and radius of PSR J0030+0451 described earlier in this Letter, for the two parameterizations of the EoS described in the text. We define the equivalent number density as the rest mass density divided by the mass of a neutron. Top row: constraints on the spectral EoS described in the text. Top left: the vertical separation of the black dashed–dotted lines shows the full range of the assumed prior for the pressure at each equivalent number density, for the spectral EoS; the black solid lines show the 5% (bottom) and 95% (top) credibility contours for the pressure at each number density, when only the priors are used. Top middle: the black dotted lines show the 5% and 95% contours in the credible region when only the prior on the pressure is used, whereas the red shaded region, which is bounded by the red dashed lines, shows the range of pressures between the 5% and 95% contours when the priors are augmented by the mass and radius estimates that we report here; the solid black line shows the median pressure within the credible region at each density. Top right: the red dashed lines are the same as in the middle panel, whereas the blue dashed lines show the soft, intermediate, and hard EoS described in Hebeler et al. (2013); the solid black lines show the 5% and 95% contours when the information provided by the measured masses of PSR J1614–2230 (Arzoumanian et al. 2018a), PSR J0348+0432 (Antoniadis et al. 2013), and PSR J0740+662 (Cromartie et al. 2019), and the tidal deformability estimate from the GW170817 event is added to the information provided by the mass and radius estimates reported here. Bottom row: the same as in the top row, but for the piecewise-polytropic EoS described in the text. These results show that the measurements of the mass and radius of PSR J0030+0451 made using *NICER* definitely tighten the constraints on the EoS of the cold, catalyzed matter in the interiors of neutron stars.

probability distributions rather than strict bounds on the mass or other quantities. They also emphasize that mass–radius measurements such as the ones we report in Section 5.3 need to be incorporated in the constraints on the EoS by fully marginalizing over the central density, rather than (for example) simply picking the point in a mass–radius curve from an EoS that has the highest likelihood from the mass–radius measurement.

There is at present no agreement on the EoS of neutron star matter at the densities expected near the centers of the most massive stars, or even on how best to incorporate the information that is currently available from nuclear and particle physics. A variety of parameterized EoS models are therefore currently being considered (for recent examples, see Tews et al. 2018; Baym et al. 2019; Greif et al. 2019). For this reason, we prefer here to use generic parameterizations to illustrate the new constraints on EoS models that are provided by the new measurements of the mass and radius of PSR J0030+0451 reported above. We caution, though, that as has been emphasized by Greif et al. (2019) and others, the constraints on the EoS derived from a given set of observations depend somewhat on the parameterization chosen for the EoS. Here we present results for the following two parameterizations.

1. The spectral parameterization of Lindblom (2010; see also Lindblom 2018). In this approach, the adiabatic index $\Gamma(p) = [(\rho + p)/p] (dp/d\rho)$ is represented using

$$\Gamma(p) = \exp\left(\sum_k \gamma_k x^k\right), \quad (4)$$

where $x \equiv \log(p/p_0)$ and p_0 is the pressure at $\rho_s/2$, where ρ_s is the density of nuclear saturation, which we take to be $\approx 2.7 \times 10^{14} \text{ g cm}^{-3}$. Following Carney et al. (2018) and Abbott et al. (2018), we carry out the expansion up to x^3 . The intervals in x that we use and the corresponding uniform priors are $\gamma_0 \in [0.2, 2]$, $\gamma_1 \in [-1.6, 1.7]$, $\gamma_2 \in [-0.6, 0.6]$, and $\gamma_3 \in [-0.02, 0.02]$.

2. A piecewise polytrope with transition densities that are also parameters. Here the adiabatic index from $\rho_s/2$ to $\rho_2 \in [3/4, 5/4]\rho_s$ is $\Gamma_1 \in [2, 3]$; from ρ_2 to $\rho_3 \in [3/2, 5/2]\rho_s$ is $\Gamma_2 \in [0, 5]$; from ρ_3 to $\rho_4 \in [3, 5]\rho_s$ is $\Gamma_3 \in [0, 5]$; from ρ_4 to $\rho_5 \in [6, 10]\rho_s$ is $\Gamma_4 \in [0, 5]$; and from ρ_5 to ∞ is $\Gamma_5 \in [0, 5]$ (all priors are uniform in the listed range). The relatively restricted range of Γ_1 is based on the study of Hebeler et al. (2013). Note also that the large number of parameters in this parameterization is intended to give flexibility to our description of the equation of state; we are not attempting to produce a concise fit to realistic microphysics. Unlike the spectral parameterization, the piecewise polytrope can represent phase transitions, in regions where the adiabatic index is close to zero.

For both parameterizations, we use the EoS of Douchin & Haensel (2001) up to $\rho = \rho_s/2$. We also use as a prior the requirement that all candidate EoS be able to support stars with masses of at least $1.4 M_\odot$. Unlike in Miller et al. (2019), we do not incorporate information about the nuclear symmetry energy, but otherwise we follow their procedure.

For a given EoS, we do not consider densities above the density at which the adiabatic sound speed for that particular

EoS exceeds the speed of light ($dp/d\rho > c^2$). We generate 150,000 random equations of state from each parameterization and then weight them, first, based only on the priors, then using our mass and radius measurement for PSR J0030+0451, and finally also including the high measured masses of PSR J1614–2230 (Arzoumanian et al. 2018a), PSR J0348+0432 (Antoniadis et al. 2013), and PSR J0740+662 (Cromartie et al. 2019), and the tidal deformability constraints from the GW170817 binary neutron star coalescence event (Abbott et al. 2018; De et al. 2018).

In order to perform our analysis, we need a way to turn discrete samples of the (M, R) posterior into an estimate of the continuous posterior probability distribution. We use the standard approach of kernel density estimation (see Rosenblatt 1956; Parzen 1962; Silverman 1986 for details). In kernel density estimation, each point in the sample is replaced by an extended probability distribution (that is, a kernel), and then the probability density at any (M, R) is proportional to the sum of the kernels associated with every point, evaluated at (M, R) . One must make a choice for both the shape of the kernel and its width (called its bandwidth in this context). For our kernels we use a Gaussian form, and a bandwidth matrix \mathbf{H} based on the covariance matrix. More specifically, our estimate of the probability density at a given point $\mathbf{x} = (M, R)$ given \mathbf{H} and samples at the points $\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_n$ is

$$\hat{f}_{\mathbf{H}}(\mathbf{x}) = \frac{1}{n} \sum_{i=1}^n \frac{1}{\det(\mathbf{H})} \mathcal{K}\{\mathbf{H}^{-1}(\mathbf{x} - \mathbf{X}_i)\}. \quad (5)$$

Here $\mathcal{K}(\mathbf{v})$, where \mathbf{v} is a column vector, is

$$\mathcal{K}(\mathbf{v}) = \exp(-\mathbf{v}^T \mathbf{v} / 2). \quad (6)$$

For \mathbf{H} we employ the covariance matrix Σ . To calculate Σ , assume that each sample point in (M, R) is a 2D vector (M_i, R_i) . Let μ_M be the mean of M over the sample, let μ_R similarly be the mean of R over the sample, and let $E[(Y_i - \mu_i)(Y_j - \mu_j)]$ be the mean of $(Y_i - \mu_i)(Y_j - \mu_j)$ over the sample where i and j are M or R . Then the ij element of Σ is

$$\Sigma_{ij} = E[(Y_i - \mu_i)(Y_j - \mu_j)]. \quad (7)$$

The standard rule-of-thumb from Silverman (1986) is that the bandwidth matrix is

$$\mathbf{H} = n^{-1/(d+4)} \Sigma^{1/2} \quad (8)$$

where d is the number of dimensions in the data ($d = 2$ in our case) and n is the number of samples from the posterior. This bandwidth is a compromise: if the width is too great then information is lost (an infinitely broad kernel would lead to a constant probability density) whereas if the width is too narrow then the probability density is choppy, with many peaks near the points that happened to be sampled.

We find, however, that our measurement of M and R for PSR J0030+0451 is precise enough that the standard bandwidth produces an inappropriately broadened posterior. For example, using the standard bandwidth we find a smoothed posterior in M/R that is $\sim 50\%$ broader than the unsmoothed posterior. We therefore multiply the standard bandwidth by 0.1, so that for n samples and two dimensions ($d = 2$),

$$\mathbf{H} = 0.1n^{-1/6} \Sigma^{1/2}. \quad (9)$$

This modification yields a smoothed posterior that retains the information without introducing choppiness.

As Figure 14 shows, the measurements of the mass and radius of PSR J0030+0451 made using *NICER* have definitely tightened the constraints on the EoS of the cold, catalyzed matter in the interiors of neutron stars.

7. Conclusions

Our analysis of the *NICER* pulse waveform data on PSR J0030+0451 and the independent analysis by Riley et al. (2019) are the first analyses of this type on data with enough counts to provide a precise and reliable measurement of the mass and radius of a neutron star. The radius and mass estimates given by our analysis are $R_e = 13.02_{-1.06}^{+1.24}$ km and $M = 1.44_{-0.14}^{+0.15} M_{\odot}$ (68%). These estimates imply an allowed range of high-density equations of state that is consistent with previous measurements of neutron star masses (see, e.g., Antoniadis et al. 2013; Arzoumanian et al. 2018a; Cromartie et al. 2019), with the tidal deformability of the neutron stars in GW170817 (see, e.g., Abbott et al. 2018; De et al. 2018), and with general nuclear physics considerations (see, e.g., the discussion in De et al. 2018). The *NICER* measurements have significantly tightened constraints on the EoS. The consistency of our results with those of Riley et al. (2019) adds confidence that the systematic errors in these results are not large, but further pulse waveform modeling and analysis of additional *NICER* data will be helpful.

In summary, our results and those of Riley et al. (2019) provide a new constraint on the radius and mass of PSR J0030+0451, and through them, on the properties of the dense matter in its core (see also Raaijmakers et al. 2019). This work, and future analyses of *NICER* data on other sources, represents an important contribution to the ongoing transition to an era of unprecedented precision in our knowledge of neutron star interiors.

This work was supported by NASA through the *NICER* mission and the Astrophysics Explorers Program. The authors acknowledge the University of Maryland supercomputing resources (<http://hpcc.umd.edu>) that were made available for conducting the research reported in this Letter. M.C.M. is grateful for the hospitality of the Kavli Institute for Theoretical Physics at the University of California, Santa Barbara, where part of this Letter was written. This Letter was therefore supported in part by the National Science Foundation under grant No. NSF PHY-1748958. M.C.M. was also supported by a Visiting Researcher position at Perimeter Institute for Theoretical Physics in the last stages of this project. W.C.G.H. appreciates the use of computer facilities at the Kavli Institute for Particle Astrophysics and Cosmology. J.M.L. acknowledges support from NASA through grant 80NSSC17K0554 and the U.S. DOE from grant DE-FG02-87ER40317. R.M.L. acknowledges the support of NASA through Hubble Fellowship Program grant *HST*-HF2-51440.001. The authors acknowledge the use of NASA's Astrophysics Data System (ADS) Bibliographic Services and the arXiv.

Facility: *NICER* (Gendreau et al. 2016).

Software: emcee (Foreman-Mackey et al. 2013), MultiNest (Feroz et al. 2009), Python and NumPy (Oliphant 2007), Matplotlib (Hunter 2007), Cython (Behnel et al. 2011), Super-Mongo (<https://www.astro.princeton.edu/~rhl/sm/sm.html>).

Appendix A Posterior Distributions for the Model With Two Oval Spots

Figure 15 shows the posterior probability density distributions for each of the parameters in the model pulse waveform produced by two, uniform, oval hot spots that best fits the NICER pulse waveform data on J0030+0451.




Figure 15. Posterior probability density distributions for the parameters in the best-fitting model pulse waveform produced by two, uniform, oval hot spots.

Appendix B

Posterior Distributions for the Model With Three Oval Spots

Figure 16 shows the posterior probability density distributions for each of the parameters in the model pulse waveform produced by three, uniform, oval hot spots that best fits the NICER pulse waveform data on J0030+0451.




















Figure 16. Posterior probability density distributions for the parameters in the best-fitting model pulse waveform produced by three, uniform, oval hot spots.

ORCID iDs

M. C. Miller  <https://orcid.org/0000-0002-2666-728X>
 F. K. Lamb  <https://orcid.org/0000-0002-3862-7402>
 A. J. Dittmann  <https://orcid.org/0000-0001-6157-6722>
 S. Bogdanov  <https://orcid.org/0000-0002-9870-2742>
 S. Guillot  <https://orcid.org/0000-0002-6449-106X>
 A. K. Harding  <https://orcid.org/0000-0001-6119-859X>
 W. C. G. Ho  <https://orcid.org/0000-0002-6089-6836>
 J. M. Lattimer  <https://orcid.org/0000-0002-5907-4552>
 R. M. Ludlam  <https://orcid.org/0000-0002-8961-939X>
 S. Mahmoodifar  <https://orcid.org/0000-0003-2386-1359>
 S. M. Morsink  <https://orcid.org/0000-0003-4357-0575>
 P. S. Ray  <https://orcid.org/0000-0002-5297-5278>
 T. E. Strohmayer  <https://orcid.org/0000-0001-7681-5845>
 K. S. Wood  <https://orcid.org/0000-0002-7376-3151>
 T. Enoto  <https://orcid.org/0000-0003-1244-3100>
 R. Foster  <https://orcid.org/0000-0002-2731-9295>

References

- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017, *PhRvL*, **119**, 161101
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2018, *PhRvL*, **121**, 161101
 Agathos, M., Meidam, J., Del Pozzo, W., et al. 2015, *PhRvD*, **92**, 023012
 Akmal, A., Pandharipande, V. R., & Ravenhall, D. G. 1998, *PhRvC*, **58**, 1804
 Alcock, C., & Illarionov, A. 1980, *ApJ*, **235**, 534
 AlGendy, M., & Morsink, S. M. 2014, *ApJ*, **791**, 78
 Alvarez-Castillo, D., Ayriyan, A., Benic, S., et al. 2016, *EPJA*, **52**, 69
 Antoniadis, J., Freire, P. C. C., Wex, N., et al. 2013, *Sci*, **340**, 448
 Arons, J. 1981, *ApJ*, **248**, 1099
 Arzoumanian, Z., Baker, P. T., Brazier, A., et al. 2018a, *ApJ*, **859**, 47
 Arzoumanian, Z., Brazier, A., Burke-Spolaor, S., et al. 2018b, *ApJS*, **235**, 37
 Badnell, N. R., Bautista, M. A., Butler, K., et al. 2005, *MNRAS*, **360**, 458
 Baldo, M., Bombaci, I., & Burgio, G. F. 1997, *A&A*, **328**, 274
 Baym, G., Furusawa, S., Hatsuda, T., Kojo, T., & Togashi, H. 2019, *ApJ*, **885**, 42
 Behnel, S., Bradshaw, R., Citro, C., et al. 2011, *CSE*, **13**, 31
 Bogdanov, S. 2013, *ApJ*, **762**, 96
 Bogdanov, S., & Grindlay, J. E. 2009, *ApJ*, **703**, 1557
 Bogdanov, S., Grindlay, J. E., & Rybicki, G. B. 2008, *ApJ*, **689**, 407
 Bogdanov, S., Guillot, S., Ray, P. S., et al. 2019a, *ApJL*, **887**, L25
 Bogdanov, S., Lamb, F. K., Mahmoodifar, S., et al. 2019b, *ApJL*, **L26**
 Bogdanov, S., Rybicki, G. B., & Grindlay, J. E. 2007, *ApJ*, **670**, 668
 Brambilla, G., Kalapotharakos, C., Timokhin, A. N., Harding, A. K., & Kazanas, D. 2018, *ApJ*, **858**, 81
 Cadeau, C., Morsink, S. M., Leahy, D., & Campbell, S. S. 2007, *ApJ*, **654**, 458
 Carney, M. F., Wade, L. E., & Irwin, B. S. 2018, *PhRvD*, **98**, 063004
 Chang, S., & Zhang, L. 2019, *MNRAS*, **483**, 1796
 Colgan, J., Kilcrease, D. P., Magee, N. H., et al. 2016, *ApJ*, **817**, 116
 Contopoulos, I., & Spitkovsky, A. 2006, *ApJ*, **643**, 1139
 Cromartie, H. T., Fonseca, E., Ransom, S. M., et al. 2019, arXiv:1904.06759
 De, S., Finstad, D., Lattimer, J. M., et al. 2018, *PhRvL*, **121**, 091102
 Demorest, P. B., Pennucci, T., Ransom, S. M., Roberts, M. S. E., & Hessels, J. W. T. 2010, *Nature*, **467**, 1081
 Douchin, F., & Haensel, P. 2001, *A&A*, **380**, 151
 Feroz, F., Hobson, M. P., & Bridges, M. 2009, *MNRAS*, **398**, 1601
 Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, **125**, 306
 Gendreau, K. C., Arzoumanian, Z., Adkins, P. W., et al. 2016, *Proc. SPIE*, **9905**, 99051H
 Greif, S. K., Raaijmakers, G., Hebel, K., Schwenk, A., & Watts, A. L. 2019, *MNRAS*, **485**, 5363
 Harding, A. K., & Muslimov, A. G. 2001, *ApJ*, **556**, 987
 Harding, A. K., & Muslimov, A. G. 2002, *ApJ*, **568**, 862
 Hebel, K., Lattimer, J. M., Pethick, C. J., & Schwenk, A. 2013, *ApJ*, **773**, 11
 Ho, W. C. G., & Heinke, C. O. 2009, *Nature*, **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
 Hunter, J. D. 2007, *CSE*, **9**, 90
 Iglesias, C. A., & Rogers, F. J. 1996, *ApJ*, **464**, 943
 Johnson, T. J., Venter, C., Harding, A. K., et al. 2014, *ApJS*, **213**, 6
 Kalapotharakos, C., Brambilla, G., Timokhin, A., Harding, A. K., & Kazanas, D. 2018, *ApJ*, **857**, 44
 Kalapotharakos, C., Harding, A. K., & Kazanas, D. 2014, *ApJ*, **793**, 97
 Lackey, B. C., & Wade, L. 2015, *PhRvD*, **91**, 043002
 Lai, D. 2001, *RvMP*, **73**, 629
 Lamb, F. K., Boutloukos, S., Van Wassenhove, S. O., et al. 2009a, *ApJ*, **706**, 417
 Lamb, F. K., Boutloukos, S., Van Wassenhove, S. O., et al. 2009b, *ApJL*, **705**, L36
 Lattimer, J. M., & Prakash, M. 2007, *PhR*, **442**, 109
 Lindblom, L. 2010, *PhRvD*, **82**, 103011
 Lindblom, L. 2018, *PhRvD*, **97**, 123019
 Lo, K. H., Miller, M. C., Bhattacharyya, S., & Lamb, F. K. 2013, *ApJ*, **776**, 19
 Lockhart, W., Gralla, S. E., Özel, F., & Psaltis, D. 2019, *MNRAS*, **490**, 1774
 Ludlam, R. M., Miller, J. M., Arzoumanian, Z., et al. 2018, *ApJL*, **858**, L5
 Miller, M. C. 2013, arXiv:1312.0029
 Miller, M. C., Chirenti, C., & Lamb, F. K. 2019, arXiv:1904.08907
 Miller, M. C., & Lamb, F. K. 2015, *ApJ*, **808**, 31
 Miller, M. C., & Lamb, F. K. 2016, *EPJA*, **52**, 63
 Morsink, S. M., Leahy, D. A., Cadeau, C., & Braga, J. 2007, *ApJ*, **663**, 1244
 Nättilä, J., Miller, M. C., Steiner, A. W., et al. 2017, *A&A*, **608**, A31
 Oliphant, T. E. 2007, *CSE*, **9**, 10
 Oppenheimer, J. R., & Volkoff, G. M. 1939, *PhRv*, **55**, 374
 Özel, F., Psaltis, D., Güver, T., et al. 2016, *ApJ*, **820**, 28
 Parzen, E. 1962, *Ann. Math. Statist.*, **33**, 1065
 Pavlov, G. G., & Zavlin, V. E. 1997, *ApJL*, **490**, L91
 Pearson, K. 1900, *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, **50**, 157
 Philippov, A. A., & Spitkovsky, A. 2018, *ApJ*, **855**, 94
 Philippov, A. A., Spitkovsky, A., & Cerutti, B. 2015, *ApJL*, **801**, L19
 Potekhin, A. Y. 2014, *PhyU*, **57**, 735
 Psaltis, D., & Lamb, F. K. 1999, *A&AT*, **18**, 447
 Raaijmakers, G., Riley, T. E., Watts, A. L., et al. 2019, *ApJL*, **887**, L22
 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
 Rosenblatt, M. 1956, *Ann. Math. Statist.*, **27**, 832
 Ruderman, M. 1991, *ApJ*, **366**, 261
 Ruderman, M., Zhu, T., & Chen, K. 1998, *ApJ*, **492**, 267
 Salmi, T., Nättilä, J., & Poutanen, J. 2018, *A&A*, **618**, A161
 Silverman, B. W. 1986, *Density Estimation for Statistics and Data Analysis* (London, Boca Raton, FL: Chapman and Hall, CRC), <http://engineering.purdue.edu/mark/pthesis/>
 Steiner, A. W., Heinke, C. O., Bogdanov, S., et al. 2018, *MNRAS*, **476**, 421
 Steiner, A. W., Lattimer, J. M., & Brown, E. F. 2010, *ApJ*, **722**, 33
 Tews, I., Carlson, J., Gandolfi, S., & Reddy, S. 2018, *ApJ*, **860**, 149
 Timokhin, A. N., & Arons, J. 2013, *MNRAS*, **429**, 20
 Tolman, R. C. 1939, *PhRv*, **55**, 364
 Tsai, Y.-S. 1974, *RvMP*, **46**, 815
 Zavlin, V. E., & Pavlov, G. G. 1998, *A&A*, **329**, 583