



Constraining the Dense Matter Equation of State with Joint Analysis of NICER and LIGO/Virgo Measurements

G. Raaijmakers¹, S. K. Greif^{2,3}, T. E. Riley⁴, T. Hinderer^{1,5}, K. Hebeler^{2,3}, A. Schwenk^{2,3,6}, A. L. Watts⁴,
S. Nisanke^{1,7}, S. Guillot^{8,9}, J. M. Lattimer¹⁰, and R. M. Ludlam¹¹

¹ GRAPPA, Anton Pannekoek Institute for Astronomy and Institute of High-Energy Physics, University of Amsterdam, Science Park 904, 1098 XH Amsterdam, The Netherlands; G.Raaijmakers@uva.nl

² Institut für Kernphysik, Technische Universität Darmstadt, 64289 Darmstadt, Germany

³ ExtreMe Matter Institute EMMI, GSI Helmholtzzentrum für Schwerionenforschung GmbH, 64291 Darmstadt, Germany

⁴ Anton Pannekoek Institute for Astronomy, University of Amsterdam, Science Park 904, 1090GE Amsterdam, The Netherlands

⁵ Delta Institute for Theoretical Physics, Science Park 904, 1090 GL Amsterdam, The Netherlands

⁶ Max-Planck-Institut für Kernphysik, Saupfercheckweg 1, 69117 Heidelberg, Germany

⁷ Nikhef, Science Park 105, 1098 XG Amsterdam, The Netherlands

⁸ IRAP, CNRS, 9 avenue du Colonel Roche, BP 44346, F-31028 Toulouse Cedex 4, France

⁹ Université de Toulouse, CNES, UPS-OMP, F-31028 Toulouse, France

¹⁰ Department of Physics and Astronomy, Stony Brook University, Stony Brook, NY 11794-3800, USA

¹¹ Cahill Center for Astronomy and Astrophysics, California Institute of Technology, Pasadena, CA 91125, USA

Received 2019 December 24; revised 2020 March 17; accepted 2020 March 23; published 2020 April 13

Abstract

The Neutron Star Interior Composition Explorer collaboration recently published a joint estimate of the mass and the radius of PSR J0030+0451, derived via X-ray pulse-profile modeling. Raaijmakers et al. explored the implications of this measurement for the dense matter equation of state (EOS) using two parameterizations of the high-density EOS: a piecewise-polytropic model, and a model based on the speed of sound in neutron stars (NSs). In this work we obtain further constraints on the EOS following this approach, but we also include information about the tidal deformability of NSs from the gravitational wave signal of the compact binary merger GW170817. We compare the constraints on the EOS to those set by the recent measurement of a $2.14 M_{\odot}$ pulsar, included as a likelihood function approximated by a Gaussian, and find a small increase in information gain. To show the flexibility of our method, we also explore the possibility that GW170817 was a NS–black hole merger, which yields weaker constraints on the EOS.

Unified Astronomy Thesaurus concepts: Neutron star cores (1107); Gravitational waves (678); Millisecond pulsars (1062); Rotation powered pulsars (1408); Bayesian statistics (1900); Nuclear physics (2077); Nuclear astrophysics (1129); X-ray astronomy (1810)

1. Introduction

Determining the behavior of matter at supranuclear densities is one of the major challenges of modern astrophysics and nuclear physics. Astronomical multi-messenger observations yield statistical measurements of neutron star (NS) properties such as gravitational mass, radius, and tidal deformability, providing a way to study matter under extreme conditions. Theoretical predictions for the phases of matter in NS cores span a wide range, from neutron-rich nucleonic matter to hyperon or deconfined quark formation or the emergence of a Bose–Einstein condensate or a color superconducting phase (see Hebeler et al. 2015; Lattimer & Prakash 2016; Oertel et al. 2017; Baym et al. 2018, for recent reviews on this topic). In practice, our uncertainty about dense matter is usually expressed in terms of a space of viable equation of state (EOS) models (see, e.g., Abbott et al. 2018; Raaijmakers et al. 2019, and references therein).

Recently NASA’s Neutron Star Interior Composition Explorer (NICER), an X-ray telescope on board the International Space Station, has delivered a joint mass–radius measurement for the millisecond pulsar PSR J0030+0451 using pulse-profile modeling (see Watts 2019 and references therein for a description of the technique). Two independent analyses were conducted within the collaboration, each making slightly different assumptions about the modeling (including priors; Miller et al. 2019; Riley et al. 2019a). The results

depended strongly on the assumed geometry for the X-ray-emitting surface hot regions, but it was possible to identify a superior configuration based principally on the likelihood. The results of the two analyses were, however, deemed consistent: Riley et al. (2019a) reported an inferred mass and equatorial radius of $M = 1.34^{+0.15}_{-0.16} M_{\odot}$ and $R_{\text{eq}} = 12.71^{+1.14}_{-1.19}$ km (for the 68% credible interval); Miller et al. (2019), on the other hand, found $M = 1.44^{+0.15}_{-0.14} M_{\odot}$ and $R_{\text{eq}} = 13.02^{+1.24}_{-1.06}$ km.

Constraints on the mass and radius have recently also been obtained from the gravitational wave (GW) observations of the binary NS merger event GW170817 (Abbott et al. 2017). The LIGO Scientific and Virgo Collaborations (LVC) reported measurements of the masses and EOS-dependent tidal deformability parameters of the NSs under different prior assumptions on the spins and using various waveform models (Abbott et al. 2017, 2019a, 2019b); see Kastaun & Ohme (2019) for a critical re-examination of the results. The corresponding radii and EOS constraints were inferred in two ways, by using a parameterized spectral EOS (Lindblom 2010) and by employing EOS-insensitive relations, both using the low-spin priors ($cS/GM^2 < 0.05$ where S is the spin angular momentum) and hence non-rotating stellar models. The results were consistent, leading to $R = 11.9^{+1.4}_{-1.4}$ km from the spectral EOS analysis and $R = 10.8^{+2.0}_{-1.7}$ km for the more massive NS at the 90% credible interval (Abbott et al. 2018); see also De et al. (2018), Annala et al. (2018), and Tews et al. (2018) for

independent related work. A number of studies have further included additional constraints from the electromagnetic counterparts assuming a NS–NS progenitor (Bauswein et al. 2017; Gao et al. 2017; Coughlin et al. 2018; Most et al. 2018; Capano et al. 2019; Margalit & Metzger 2019; Radice & Dai 2019; Shibata et al. 2019). In this work, we remain cautious of the large uncertainties in modeling the electromagnetic counterparts and only use the fact that the observed kilonova, an ultraviolet–optical–infrared transient powered by rapid neutron-capture nucleosynthesis (see, e.g., Lattimer & Schramm 1976; Li & Paczyński 1998; Rosswog et al. 1999; Kulkarni 2005; Metzger et al. 2010; Metzger 2017) indicated that the progenitor binary involved at least one NS. We consider two possibilities, a double NS system as assumed in most analyses and a NS–black hole (BH) binary. The mass of the BH in the latter scenario would be very low and could have originated from an earlier merger of two NSs or be of primordial origin (see, e.g., Yang et al. 2018; Coughlin & Dietrich 2019; Hinderer et al. 2019).

In this Letter, we perform a joint analysis of the EOS constraints from NICER and GW170817, following the method for connecting global NS parameters to the EOS used in Raaijmakers et al. (2019). We focus on the mass–radius measurement from Riley et al. (2019a),¹² as the measurement of Miller et al. (2019) has already been used to jointly constrain the EOS with GW and radio pulsar measurements. We will compare our findings to those of Miller et al. (2019) in Section 4.

We use EOS models that incorporate prior information from nuclear physics up to around saturation density, and two different parameterized extensions at high density; one using piecewise-polytropes, and one based on physically motivated assumptions about the speed of sound. We develop the methodology for the combined interpretation of these measurements in a Bayesian framework that also takes into account the measurements of massive pulsars. Our method can readily include a larger number of NSs from anticipated future multi-messenger observations. Next, we analyze the impact of systematic uncertainties arising from different priors for the EOS. We show that the priors used for the spectral EOS inference from GWs (Abbott et al. 2018) allow for much stiffer EOSs than the priors in Raaijmakers et al. (2019) and explain the reasons for these differences. Nevertheless, we find that the resulting EOS constraints are broadly consistent. We quantify explicitly that the fact that GW measurements determine the chirp mass to high accuracy can be utilized to accelerate the parameter inference by treating it as fixed.

2. Inference Framework

We adopt the framework outlined previously in Raaijmakers et al. (2019) and Greif et al. (2019), which we will briefly summarize here, including details of some adjustments made to incorporate information from the GW data of GW170817. Note that the Bayesian methodology is very similar to that outlined in Miller et al. (2020), although there are differences in the prior assumptions (see Section 4).

¹² The posterior samples from that analysis are available in a Zenodo repository (Riley et al. 2019b). We have now updated the repository from the version published alongside Riley et al. (2019a) to include the following additional material: files containing only the mass–radius samples, contour files for the credible regions in mass–radius space, and simplified coordinate files for the boundaries of the hot regions on the stellar surface.

2.1. Parameterizations

Two distinct parameterizations are considered: a three-piece polytropic (PP) model with varying transition densities between the polytropes (Hebeler et al. 2013), and a speed of sound (CS) model based on physical considerations both at nuclear and high densities (Greif et al. 2019). Both models are matched to an additional polytrope below $1.1 n_0$ (with saturation density $n_0 = 0.16 \text{ fm}^{-3}$) with varying normalization that captures the range of allowed EOS calculated from chiral effective field theory (chiral EFT) interactions (Hebeler & Schwenk 2010; Hebeler et al. 2013). At densities below $0.5 n_0$ this polytrope is matched to the BPS crust EOS (Baym et al. 1971).

2.2. Bayesian Parameter Estimation

We use Bayesian methodology to estimate parameters in our EOS model. A more in-depth discussion on parameter estimation frameworks in the context of dense matter inference can be found in Riley et al. (2018), which we will very briefly describe here. Let us write, using Bayes’ theorem, the posterior distributions of the EOS parameters θ and central energy densities ε (together the *interior* parameters) as

$$p(\theta, \varepsilon | \mathbf{d}, \mathcal{M}) \propto p(\theta | \mathcal{M}) p(\varepsilon | \theta, \mathcal{M}) p(\mathbf{d} | \theta, \mathcal{M}), \quad (1)$$

where \mathcal{M} is the model that includes all assumed physics, and \mathbf{d} is the data set from both NICER observations and strain data of GW170817 from the GW detectors LIGO/Virgo. Note that the prior on the central energy densities ε is dependent on the EOS parameters θ , because the maximum stable central energy density is different for each set of θ . Given that the two observations are independent, we can separate the likelihood function as¹³

$$\begin{aligned} p(\theta, \varepsilon | \mathbf{d}, \mathcal{M}) &\propto p(\theta | \mathcal{M}) p(\varepsilon | \theta, \mathcal{M}) \\ &\times p(\Lambda_1, \Lambda_2, M_1, M_2 | \mathbf{d}_{\text{GW}}) \\ &\times p(M_3, R_3 | \mathbf{d}_{\text{NICER}}) \end{aligned} \quad (2)$$

where the nuisance-marginalized likelihood functions of (i) M_3 and R_3 , and (ii) $\Lambda_1, \Lambda_2, M_1$, and M_2 , are equated¹⁴ to the nuisance-marginalized joint posterior density distributions inferred by Riley et al. (2019a) and Abbott et al. (2019b), respectively. The marginal GW likelihood function is degenerate under exchange of binary components, but we adopt the same convention as in Abbott et al. (2019a) and define $M_1 \geq M_2$. The interior parameters θ and ε (now containing three central energy densities, corresponding to one observed star by NICER and two by the LVC) map deterministically to the parameters M, R , and Λ (where we have assumed the rotation of the star, Ω , to be zero¹⁵),

¹³ Note that for notational simplicity, we omit conditional arguments that would denote the model used by a collaboration. The global model \mathcal{M} can be considered as a proper superset of the union of these models.

¹⁴ For NICER, the nuisance-marginalized likelihood function $p(\mathbf{d}_{\text{NICER}} | M, R) \propto p(M, R | \mathbf{d}_{\text{NICER}})$ because the joint prior $p(M, R)$ was flat (Riley et al. 2019a). For GW170817 the nuisance-marginalized likelihood function $p(\mathbf{d}_{\text{GW}} | M_1, M_2, \Lambda_1, \Lambda_2) \propto p(M_1, M_2, \Lambda_1, \Lambda_2 | \mathbf{d}_{\text{GW}})$ because the priors used in Abbott et al. (2019a) are flat in both masses and tidal deformabilities.

¹⁵ See Raaijmakers et al. (2019) for a discussion on the spin of PSR J0030+0451. We assume that the spins of the two components in GW170817 are consistent with the spins of observed Galactic binary NS systems that will merge within a Hubble time. Therefore we use the posterior distribution of GW170817 in the case of a low-spin prior, i.e., $\chi < 0.05$, which corresponds to slower rotation than PSR J0030+0451.

allowing us to sample from the prior distribution of θ and ε and then numerically evaluate the likelihood functions using kernel density estimation (KDE) on the posterior samples. We then draw from the joint posterior distribution $p(\theta, \varepsilon | \mathbf{d}, \mathcal{M})$ of all interior parameters.

However, one complication that arises when performing KDE on samples from the joint posterior distribution of the two masses associated with GW170817 is that, due to the extreme accuracy to which the chirp mass $M_{\text{chirp}} = (M_1 M_2)^{3/5} / (M_1 + M_2)^{1/5}$ is known relative to the uncertainty in the individual masses, the choice of bandwidth is difficult to make (for GW170817, $M_{\text{chirp}} = 1.186 \pm 0.001 M_\odot$; Abbott et al. 2019a). A small bandwidth is necessary to accurately describe the chirp mass, while a larger bandwidth is necessary to smooth out finite sampling noise in the distribution of masses. Another complication is that when the two sampled central densities are uncorrelated—except for the assumption of a shared EOS—it is computationally expensive for samplers to find the region in the space of masses where all of the probability density is concentrated.

To avoid these complications, and at the same time utilize the small chirp mass uncertainty, we fix it to its median value of $M_{\text{chirp}} = 1.186 M_\odot$. Consequently, the mass of the secondary object is a deterministic function of the mass of the primary object, and there is one fewer free central density parameter in the vector ε . For details on the approximation invoked here, we refer to Appendix A, up to Equation (A12). For likelihood evaluation we use the mass ratio $q = M_2/M_1$ instead of the individual masses, transforming Equation (2) to¹⁶

$$\begin{aligned} p(\theta, \varepsilon | \mathbf{d}, \mathcal{M}) &\propto p(\theta | \mathcal{M}) p(\varepsilon | \theta, \mathcal{M}) \\ &\times p(\Lambda_1, \Lambda_2, q | \mathbf{d}_{\text{GW}}) \\ &\times p(M_3, R_3 | \mathbf{d}_{\text{NICER}}). \end{aligned} \quad (3)$$

Moreover, the deformability $\Lambda_2 = \Lambda_2(\theta; q)$.

In Figure 1¹⁷ we compare the updated prior distribution after including information from GW170817 with and without fixing the chirp mass. When the chirp mass is considered as a free parameter we include it as an element of θ ,¹⁸ sample from a uniform prior $M_{\text{chirp}} \sim U(1.180, 1.192)$, and define the likelihood function as $p(\Lambda_1, \Lambda_2, q, M_{\text{chirp}} | \mathbf{d}_{\text{GW}})$, thus requiring four-dimensional KDE. The two distributions are, due to the small uncertainty in the chirp mass, almost equal, apart from some finite sampling noise. In the following, we fix the chirp mass to $M_{\text{chirp}} = 1.186 M_\odot$ in order to reduce the dimensionality of the parameter vector θ .

We use the nested sampling software MULTINEST to draw weighted samples from the posterior distribution of θ (Feroz & Hobson 2008; Feroz et al. 2009, 2013; Buchner et al. 2014).




Figure 1. Comparison of the prior on EOS parameters, transformed to the space of mass and radius, as it would be updated after performing parameter estimation on GW170817 with (green contours) and without (blue contours) fixing the chirp mass to the median, $M_{\text{chirp}} = 1.186 M_\odot$, of its marginal posterior distribution. The two distributions show some small-scale differences in the 1σ contour but are globally consistent. For comparison we also show the prior distribution before including information from GW170817, but with the $2.14 M_\odot$ pulsar information, in black contours. For all contours the dotted and dashed lines indicate the 68% and 95% credible regions, respectively.

2.3. Priors

The bounds of the prior ranges of parameters used in this analysis are identical to those discussed in Section 2.1.1 of Raaijmakers et al. (2019) and Section 3.1.1 of Greif et al. (2019). Within these bounds all EOS parameters are sampled uniformly, while the central density of a star, ε_c , is uniformly sampled in logarithmic space between $\log(\varepsilon_c) = 14.6$ and an upper bound that is determined by the maximum central density of each particular EOS. We consider an EOS in the PP model up until the highest density that corresponds to a stable NS or to the point where causality is no longer satisfied, i.e., where the speed of sound exceeds the speed of light, $c_s > c$. The CS model has slightly more restrictive requirements.

- (i) The speed of sound for all densities should be lower than the speed of light.
- (ii) At asymptotic densities ($\sim 50n_0$) the speed of sound should converge to $c_s = \sqrt{1/3}c$ from below, based on theoretical calculations of c_s in the framework of perturbative quantum chromodynamics (Fraga et al. 2014).
- (iii) At low densities the speed of sound can be described by that of a normal Fermi liquid such that we require $c_s \leq \sqrt{0.163}c$ at densities below $1.5 n_0$ (for more details, see Greif et al. 2019).

A notable change from the prior used in Raaijmakers et al. (2019) is how we implement information from pulsar mass measurements in our analysis.¹⁹ These high-precision measurements obtained from the timing of radio pulsars restrict softer EOSs by requiring each EOS to be able to support the heaviest NSs. There have been several massive NSs detected, with the most stringent constraints coming from PSR J0348+0432 with a mass of $2.01^{+0.04}_{-0.04} M_\odot$ (Antoniadis et al. 2013) and more recently PSR J0740+6620 with a mass of $2.14^{+0.10}_{-0.09} M_\odot$.

¹⁶ Note that, by transforming the parameters M_1 and M_2 to the mass ratio q and M_{chirp} , the LVC prior $p(q, M_{\text{chirp}})$ is no longer flat. The likelihood function we approximated as proportional to the posterior density $p(\Lambda_1, \Lambda_2, q | \mathbf{d}_{\text{GW}})$ is therefore contaminated. We have checked, however, that reweighting the GW170817 posteriors such that the prior $p(q, M_{\text{chirp}})$ is flat has an unimportant effect on the likelihood function used for EOS inference. Moreover, note that approximating the conditional distribution $p(\Lambda_1, \Lambda_2, q | \mathbf{d}_{\text{GW}}, M_{\text{chirp}})$ by the marginal distribution $p(\Lambda_1, \Lambda_2, q | \mathbf{d}_{\text{GW}})$ has a similarly unimportant effect.

¹⁷ See the Zenodo repository (Raaijmakers et al. 2020) for the data and code to recreate all plots in this Letter.

¹⁸ Thus now mixing interior parameters and exterior spacetime parameters.

¹⁹ We refer the reader to the discussion in Section 4.2 of Raaijmakers et al. (2019), and to the arguments in Section 4.1 of Miller et al. (2020).

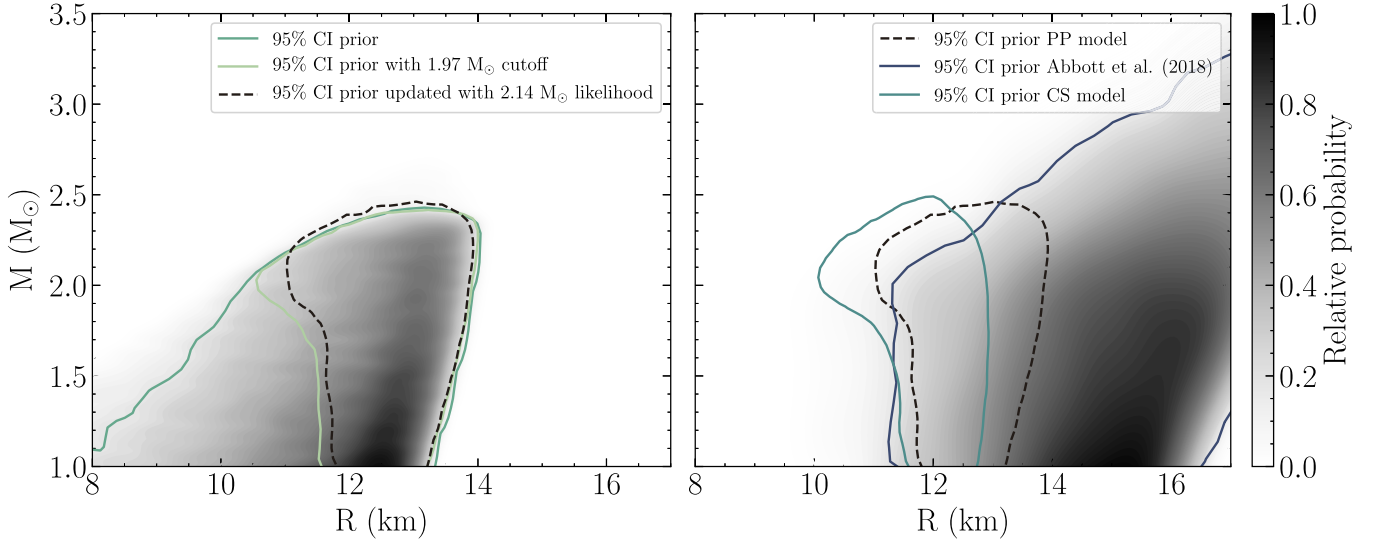


Figure 2. Left panel: comparison between the full prior distribution (black shaded region, 95% credible region bounded by the dark green contour) for the PP model with a $1.97 M_{\odot}$ cutoff (light green contour, as used in Raaijmakers et al. 2019) and when updated by parameter estimation using the $2.14 M_{\odot}$ pulsar likelihood function (black dashed contour) from Cromartie et al. (2019). Using a cutoff in the prior allows for slightly smaller radii than using the likelihood function: this is due to both the higher mass of the center of the pulsar likelihood function and the fact that the likelihood function gives more weight to an EOS with a maximum mass of $2.14 M_{\odot}$ than, e.g., $2.05 M_{\odot}$. Right panel: comparison between the 95% credible regions of the prior distributions of the PP model (black, dashed contour) and the CS model (dark green contour), when updated with the $2.14 M_{\odot}$ pulsar likelihood function. We also show the 95% credible region (black shaded region with blue contours) of the prior distribution using the spectral model (Lindblom 2018) that was used in Abbott et al. (2018). From the comparison it is clear that the prior with the spectral model allows for a much stiffer EOS but has a similar bound at low radii.

(Cromartie et al. 2019). In many previous analyses the lower 1σ limit of such a mass measurement was taken as the minimum mass that an EOS has to support a priori. However, Miller et al. (2020, 2019) do not make such an assumption and emphasize that likelihood information about high-mass pulsars be treated accurately (see also Alvarez-Castillo et al. 2016). Here we approximate the highest pulsar mass measurement as a Gaussian likelihood function,²⁰ such that Equation (3) reads

$$\begin{aligned}
 p(\theta, \varepsilon | \mathbf{d}, \mathcal{M}) &\propto p(\theta | \mathcal{M}) p(\varepsilon | \theta, \mathcal{M}) \\
 &\quad \times p(\Lambda_1, \Lambda_2, q | \mathbf{d}_{\text{GW}}) \\
 &\quad \times p(M_3, R_3 | \mathbf{d}_{\text{NICER}}) \\
 &\quad \times p(M_4 | \mathbf{d}_{\text{radio}}).
 \end{aligned} \tag{4}$$

The vector ε now contains a fourth central energy density (drawn from the same prior as described at the beginning of this section) for which the corresponding mass, M_4 is used to evaluate the Gaussian likelihood function. We compare the effect of a cutoff in the prior with the implementation of the new likelihood function for the PP model in the left panel of Figure 2. The solid lines indicating the 95% credible region show that the prior distribution between the two methods is very similar, although the likelihood implementation of the $2.14 M_{\odot}$ pulsar is slightly more constraining at lower radii due to the higher mass. In the right panel of Figure 2 we compare the prior distributions for the two parameterizations used in this Letter with the prior distribution of the spectral parameterization used in Abbott et al. (2018). The spectral parameterization allows for much larger radii than we consider here, due to using only a crust EOS without implementing nuclear physics

constraints around nuclear saturation density. This is taken into account in this work (see also Hebeler et al. 2013) by adopting the EOS band based on chiral EFT up to $1.1n_0$. The exact breakdown density of chiral EFT is not fully known, but many calculated and also predicted nuclear properties are consistent with experiment (Hebeler et al. 2015), including the symmetry energy and other matter properties at saturation density, suggesting that the range of possible EOSs predicted by chiral EFT is valid up to around nuclear saturation density (see also Section 4.2 of Raaijmakers et al. 2019).

2.4. Generalization to a Large Number of Stars

The separation of the likelihood function based on different observables in Equation (4) is a useful way of analyzing multiple sources at the same time. In the future, however, when a population of NSs is observable, this can quickly become computationally intractable—depending on the sampling algorithm applied—as the number of densities constituting the parameter vector ε grows linearly with the number of observed stars. One can perform the parameter estimation sequentially in this case, where the prior $p(\theta | \mathcal{M})$ is updated after each iteration and sampled from in the next (see also Figure 2.1 in Riley et al. 2018).²¹

3. EOS Constraints Given Multi-messenger Observations

For both parameterizations discussed in Section 2.1 we draw weighted samples from the posterior distribution $p(\theta | \mathbf{d}, \mathcal{M})$ using nested sampling. In order to explore the effect of different measurements on this posterior distribution, we start by only considering information from the $2.14 M_{\odot}$ pulsar; we then

²⁰ A value of $0.09 M_{\odot}$ was chosen for σ in the Gaussian likelihood function, which is not representative for the upper tail of the pulsar mass distribution. We believe, however, that the effect on the posterior distributions of the EOS parameters is small enough to justify this approximation.

²¹ Although natural, sequential updating does not come without its own technical challenges. Moreover, in a rigorous population-level context, one should consider the role of hyperparameters.

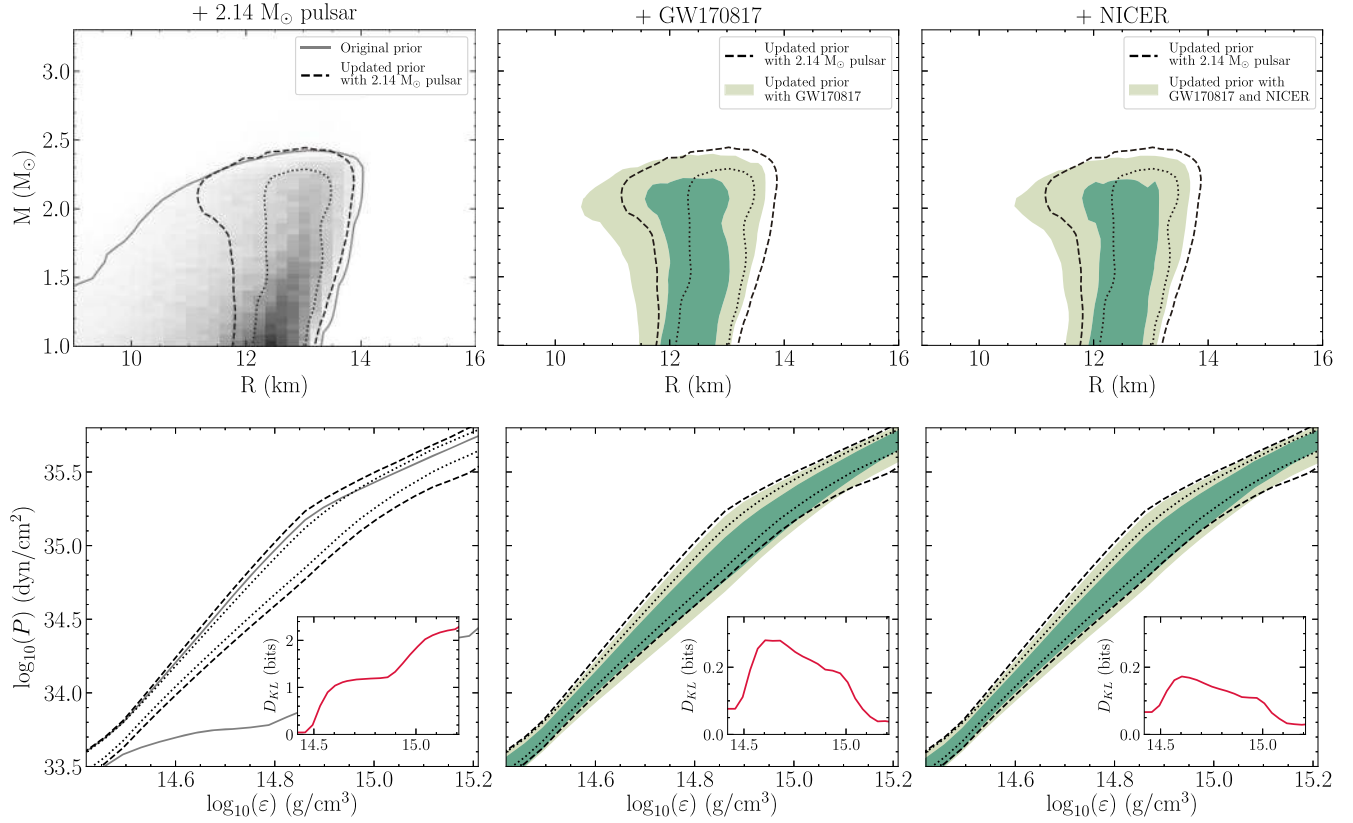


Figure 3. Posterior distributions conditional on the PP model and given: (i) the $2.14 M_{\odot}$ pulsar alone (left panels), (ii) inclusion of the GW170817 measurements (middle panels), and (iii) inclusion of the mass and radius of PSR J0030+0451 inferred by Riley et al. (2019a) given NICER data (right panels). In the top row we show how the posterior distributions update the prior distributions, by drawing a new central density given the inferred distribution on EOS parameters, $p(\varepsilon_c | \text{EOS})$. This is then transformed to the space of masses and radii, with the contours indicating the 68% and 95% credible intervals. In the bottom row we show the marginal posterior distributions of the pressure P conditional on energy density ε , i.e., $p(P | \varepsilon, \mathbf{d}, \mathcal{M})$. The bands show the connected 68% and 95% credible intervals at each energy density ε . The gray lines in the left panels show the 95% credible interval of the full prior, while the black dotted and dashed lines in all panels show the 68% and 95% credible regions of the updated prior when including information from the $2.14 M_{\odot}$ pulsar. The green contours show the same credible regions, but for posterior distributions that are inferred from multiple measurements of NS observables. In the lower right inset panels we illustrate the evolution of the Kullback–Leibler divergence as a function of energy density. We conclude that most information is gained from including the $2.14 M_{\odot}$ pulsar. The binary merger GW170817 favors softer EOSs than the prior, but the measured radius from PSR J0030+0451 favors stiffer EOSs, resulting in a final posterior distribution very similar to the prior.

include information from the binary merger GW170817, and finally we include the more recent NICER measurements of PSR J0030+0451. The GW data we use here are the publicly available posterior samples (LIGO-Virgo collaboration 2018) which assume certain priors on the GW parameters as described in Abbott et al. (2019b); a study of the impact of changing these priors is outside the scope of this work but see T. Zhao et al. (2020, in preparation).

3.1. EOS Constraints Assuming GW170817 Was a NS–NS

We illustrate the posterior distribution in two different ways in Figures 3 and 4. The lower panels show the 68% and 95% credible intervals on the pressure at each energy density given the inferred distribution on EOS parameters, i.e., $p(P | \varepsilon, \mathbf{d}, \mathcal{M})$. The upper panels show the updated prior distribution for a new star given the inferred distribution on EOS parameters, transformed to the space of mass and radii. More practically, this means that for each EOS in the posterior distribution $p(\theta | \mathbf{d}, \mathcal{M})$ we draw a new central density from the prior $p(\varepsilon | \theta)$, where we use again a uniform prior in log space with the upper bound determined by the maximum stable neutron star of that EOS. From this central energy density we

then compute the corresponding mass and radius pair. By combining all pairs for each EOS sample in the posterior we obtain a distribution of masses and radii. The contours again indicate the 68% and 95% credible regions. Visually inspecting the posterior distributions for both the PP model (Figure 3) and the CS model (Figure 4) indicates that the inferred masses and tidal deformabilities for GW170817 favor softer EOSs than our prior. Folding in information about the radius of PSR J0030+0451, however, with a peak value around 12.7 km, favors stiffer EOSs. As a result the final posterior distribution is only slightly shifted toward smaller radii but otherwise closely follows the distribution when only the highest-mass pulsar is included. We also note that one should be careful when inferring the maximum mass allowed for a NS from Figures 3 and 4 because the high-mass end of these distributions depends sensitively on the prior on central energy densities. See also Section 4.2 and Figure 6 where we explicitly show the distribution of the maximum mass by taking the highest allowed central energy density from each EOS instead of sampling from the prior on ε .

In order to quantify this we compute the Kullback–Leibler (KL) divergence (Kullback & Leibler 1951) between the two distributions shown in the lower panels of Figures 3 and 4 at a




Figure 4. Same as in Figure 3 but for the CS model. Again we conclude that constraints from GW170817 point to softer EOSs with lower radii, but the results from NICER point to stiffer EOSs with higher radii. The final posterior distribution, conditional on the three different measurements combined, is then very similar to the distribution with only information from the $2.14 M_{\odot}$ pulsar included.

given energy density ϵ (see lower right inset panels). The KL divergence is an asymmetric measure of how one probability distribution differs from another; when computed using a logarithm of base two, the divergence has units of bits. As expected, most of the information is gained from folding in the $2.14 M_{\odot}$ pulsar constraint. The posterior distribution given GW170817 alone exhibits greater divergence from the prior than does the posterior distribution given GW170817 and NICER information; that said, both divergences are small at all densities.

Finally, we compute the Bayes' factors to investigate whether one parameterization is favored over the other by the data. Assuming the two discrete models to have equal probability a priori, the Bayes' factor reduces to the ratio of the evidences of the two posteriors. We quote the values for the three different posterior distributions in Table 1, where the Bayes' factor K is the ratio of the PP model over the CS model. To interpret the values of K we follow the table of Kass & Raftery (1995) and conclude that none of the Bayes' factors shows substantial support for one of the models over the other.

3.2. EOS Constraints Assuming GW170817 Was a NS–BH

Based on the GW signal and observed electromagnetic counterpart from GW170817 there is a non-negligible chance that one of the compact objects involved in the merger was a light BH (Yang et al. 2018; Ascenzi et al. 2019; Coughlin & Dietrich 2019; Hinderer et al. 2019), provided that the objects

Table 1
Log-evidences (Z) and Bayes' Factors (K)

Posterior	log (Z)		K
	PP model	CS model	
+ $2.14 M_{\odot}$ pulsar	-1.69 ± 0.03	-2.22 ± 0.02	0.70
+ GW170817	-15.44 ± 0.02	-15.08 ± 0.02	1.70
+ NICER	-17.05 ± 0.03	-17.10 ± 0.03	1.05

Note. Log-evidences (Z) and Bayes' factors (K) for the three posterior distributions and two parameterizations. Also quoted are the Bayes' factors (K), computed as the ratio of the evidence for the PP model over that for the CS model. Following the interpretation of Kass & Raftery (1995) there is no significant support for one parameterization over the other.

had an unequal mass ratio. Such a light BH could for example be formed during an earlier merger of two NSs, or originate from density fluctuations in the early universe (García-Bellido et al. 1996).

We show that making different assumptions related to the nature of the merger can affect the inferred EOS. More specifically we investigate the impact of GW170817 being a NS–BH by fixing the tidal deformability of the heavier object Λ_1 to zero. One complication, however, is that the posterior samples provided in Abbott et al. (2019a) do not contain enough samples when we restrict $\Lambda_1 = 0$. Instead, we perform a coordinate transformation on the posterior samples to the effective tidal deformability $\hat{\Lambda}$, a combination of the two




Figure 5. Comparison between the posterior distributions obtained for the assumption that GW170817 is a NS–NS or NS–BH merger, without (left panel) and with (right panel) the NICER measurements of PSR J0030+0451. Here we only show results for the PP model. The lines show the connected 95% credible regions at a given energy density ϵ . The black lines show the same credible regions when the posterior distribution is only informed by the $2.14 M_{\odot}$ pulsar. The lower right inset illustrates the evolution of the Kullback–Leibler divergence with energy density ϵ and indicates that GW170817 is slightly more constraining for the EOS when assumed to be a NS–NS merger.

independent tidal deformabilities and masses. To incorporate the constraints on the NS–BH scenario from electromagnetic (EM) observations of the associated kilonova we follow the approach outlined in Hinderer et al. (2019). We use the model of Foucart et al. (2018) to compute the remnant mass outside the BH after the merger M_{rem} from the progenitor NS compactness, which is related to its tidal deformability Λ_2 , the dimensionless spin of the BH χ_1 , and the mass ratio. For the spin we only consider aligned or anti-aligned spins with the orbital angular momentum. We obtain these inputs from the GW data by transforming the posterior samples of GW170817 to $\tilde{\Lambda}$ and χ_{eff} , the effective dimensionless spin, setting the BH tidal deformability $\Lambda_1 = 0$ and the NS spin $\chi_2 = 0$. We apply a conservative cut to the GW170817 posterior by only considering samples that have $M_{\text{rem}} > 0.1 M_{\odot}$, which can produce the observed bolometric light curve of GW170817 (Kasliwal et al. 2017) within some error margins based on semi-analytic light curve modeling. Effectively this is a crude approximation to the likelihood function that is zero if M_{rem} is below this threshold and uniform if above. The posterior distribution of the EOS parameters is then given by

$$\begin{aligned}
 p(\theta, \epsilon | \mathcal{d}, \mathcal{M}) &\propto p(\theta | \mathcal{M}) p(\epsilon | \theta, \mathcal{M}) \\
 &\times p(\tilde{\Lambda}(\Lambda_2, M_1, M_2), q | \mathcal{d}_{\text{GW,EM}}) \\
 &\times p(M_3, R_3 | \mathcal{d}_{\text{NICER}}) \\
 &\times p(M_4 | \mathcal{d}_{\text{radio}}),
 \end{aligned} \quad (5)$$

in which the tidal deformability of the heavier object is fixed to zero and the chirp mass is again fixed to $M_{\text{chirp}} = 1.186 M_{\odot}$. Equation (5) requires that the joint prior density of $\tilde{\Lambda}$ and q be *flat*. However, the prior density is *not* flat. In particular, the prior distribution of $\tilde{\Lambda}$ exhibits undesirable behavior near zero, where prior support drops to zero. To minimize the induced bias we reweight the posterior samples to a flat prior in $\tilde{\Lambda}$ using the method described in Abbott et al. (2020). The difference is, however, small since the EM information already excludes values of $\tilde{\Lambda}$ close to zero. Finally we note that the term $p(\tilde{\Lambda}(\Lambda_2, M_1, M_2), q | \mathcal{d}_{\text{GW,EM}})$ is marginalized over Λ_1 and therefore the BH–NS likelihood

function is contaminated with the likelihood of a system with two deformable objects. However, comparing the marginalized distribution with the conditional distribution $p(\tilde{\Lambda}(\Lambda_2, M_1, M_2), q | \mathcal{d}_{\text{GW,EM}}, \Lambda_1)$ shows that the approximation is valid here in the context of illustrating the impact of different assumptions about the binary (see Appendix B).

In Figure 5 we compare the posterior distribution for the EOS for the assumption that GW170817 is a NS–NS or NS–BH merger, with the prior distribution obtained from only the $2.14 M_{\odot}$ pulsar measurement. In the left panel we only consider GW170817 and the pulsar mass, while in the right panel we also include the NICER measurement. Based on the KL divergence as a function of energy density, we conclude that in both cases GW170817 has more support for softer EOSs when assumed to be a double NS merger. This can be explained by the fact that the relatively low inferred value of $\tilde{\Lambda}$ in Abbott et al. (2019a) can be achieved with a higher value of Λ_2 (i.e., stiffer EOS) when $\Lambda_1 = 0$ and the fact that the NS–BH scenario is only consistent with the EM counterpart for unequal mass ratios and larger NSs.

4. Discussion

In this work we have analyzed the combined constraints on the dense matter EOS given the recent inferred mass and radius of PSR J0030+0451 by Riley et al. (2019a) using NICER data, and the measurement of the GW signal from GW170817, in combination with the radio measurement of a $2.14 M_{\odot}$ pulsar.

4.1. Multi-messenger Contributions

The posterior distributions show that the most information is gained from the most massive pulsar mass measurement. In combination with the restricted range of possible EOS at lower densities described by the chiral EFT band and, for the CS model, by the approximation of NS matter as a Fermi liquid, the pulsar mass already puts stringent constraints on the EOS. When including information from GW170817, softer EOSs are




Figure 6. Posterior distributions of the maximum mass of a non-spinning NS M_{\max} that is supported by the inferred EOS after including the pulsar mass measurement, GW170817 and the results from NICER. For the PP model we find a median value of $M_{\max} = 2.26^{+0.16}_{-0.24} M_{\odot}$ while for the CS model we find $M_{\max} = 2.13^{+0.26}_{-0.22} M_{\odot}$, where the upper and lower limit bound the 95% credible region.

yielded a posteriori, but this is only a small effect because EOSs consistent with the $2.14 M_{\odot}$ pulsar mass measurement are restricted by causality to higher radii. Finally, the recent NICER measurement shows more support for stiffer EOSs, causing the posterior distribution to have a narrow peak where the likelihood functions of NICER and GW170817 overlap with the information from the $2.14 M_{\odot}$ pulsar and the chiral EFT band.

In order to quantitatively assess the prior-to-posterior information gain through sequential multi-messenger updates, we have computed KL divergences as a function of energy density. The divergences indicate that most information is gained from the radio pulsar mass measurement. The information gain from GW170817, given prior radio information, is small. Further, including the NICER likelihood yields a smaller information gain because the radio and X-ray modeling constrain a similar part of the EOS parameter space, i.e., provide less support for softer EOSs.

Note that these divergences all depend on which distribution is compared with which. For example, comparing the constraints from NICER or GW170817 with the original, more diffuse prior distribution would yield a higher KL divergence. We argue, however, that a logical, unique order of precedence would be ideal. An obvious option is to chronologically incorporate the various (astronomical) measurements that constrain the EOS.

It is however difficult to design an update order that is truly chronological, given that (i) not all constraints are compiled in any one analysis, and (ii) multi-messenger constraints are being derived contemporaneously, both given newly acquired data, and archival data when the modeling procedure is revolutionized. Typically there will not be a clear chronology, and even if there were, we might try to account for the number of physical assumptions a constraint is conditional on. If we assume that the number of assumptions anti-correlates with robustness to systematic error in our models of reality, we can

attempt to compile information—and archive constraints—in loose order of robustness.²²

We therefore opt to start with information contributed from radio timing of pulsars in relativistic binaries. There are multiple constraints to consider: constraints for two systems were reported before GWs were first detected (Demorest et al. 2010; Antoniadis et al. 2013), with measurements for the very first being updated in recent years with continued timing (Fonseca et al. 2016; Arzoumanian et al. 2018; Cromartie et al. 2019; Miller et al. 2020). It is believed that radio measurements relying solely on the relativistic Shapiro delay are robust to systematic error. However, the first report of a pulsar with a mass above $2 M_{\odot}$ is dependent on theoretical models of white dwarf evolution (Antoniadis et al. 2013). The issue of choice can of course be straightforwardly nullified by incorporating all of the radio pulsar mass measurements, each of which encodes orthogonal information in a population-level context. In this work we chose to use a single astronomical source for each class of astronomical messenger. Of the two highly informative measurements relying solely on the relativistic Shapiro delay, derived by Arzoumanian et al. (2018) and Cromartie et al. (2019), we chose to use the constraint reported by the latter.

4.2. Implications for the Maximum Mass of NSs

The maximum mass of NSs places important constraints on the EOS and can aid in determining the nature of a compact binary merger and the following merger remnant. In Figure 6 we show the posterior distributions for M_{\max} for the PP and CS

²² In principle, if a chronologically earlier constraint is biased, future unbiased, informative constraints should offer a strong opinion. Thus bias should resolve in time, provided that computation can be performed accurately (Raaijmakers et al. 2019). This is especially true if models are revolutionized and used to reanalyze data, at the expense of a clear chronology. In practice, however, it is not always straightforward to accurately update knowledge with future information when resolution would be required in the tail of a prior distribution.

parameterizations considered in this Letter. These distributions are obtained by calculating for each EOS sample in the posterior distribution $p(\theta|d, \mathcal{M})$ the maximum mass for that EOS. The gray dashed lines indicate the prior support on M_{\max} , which drops down rapidly above $\sim 2.5 M_{\odot}$ for both parameterizations. The bump around $2.3 M_{\odot}$ for the PP model is a consequence of EOSs that terminate at lower maximum masses when their sound speed reaches the speed of light before reaching a maximally stable NS. This leads to a higher value for the PP model of $M_{\max} = 2.26^{+0.16}_{-0.24} M_{\odot}$ compared to the CS model, $M_{\max} = 2.13^{+0.26}_{-0.22} M_{\odot}$, when including information from pulsar mass measurements, GW170817, and NICER. Here we quote median values and upper and lower limits that bound the 95% credible region. We find that the maximum mass distribution presented here is broadly consistent with values found in other works (see, e.g., Margalit & Metzger 2017; Shibata et al. 2017; Alsing et al. 2018; Rezzolla et al. 2018; Ruiz et al. 2018; Tews & Schwenk 2020). We find, however, a lower upper-limit since some of these analyses also include information from the kilonova light curves that suggest the formation of a metastable NS and disfavor the prompt formation of a BH.

4.3. Comparison to Other Work

We first compare our results with the analysis of Miller et al. (2019), who use a similar Bayesian approach to combine results from NICER, GW170817, and radio pulsar measurements to constrain the EOS. The parameterization of the EOS used in this work differs, however, in several aspects from the two parameterizations used by Miller et al. For the latter the crust EOS of Douchin & Haensel (2001) is considered up to $0.5 \rho_s$. Beyond $0.5 \rho_s$, two different extrapolations to higher densities are used. One approach used by Miller et al. is to implement the spectral parameterization by Lindblom (2010, 2018). The second parameterization is a PP expansion with two more polytropic segments than used in this work, totaling up to five segments. The range of the first polytropic index of Miller et al. (2019) is chosen to be rather restrictive in the context of chiral EFT (Hebeler et al. 2013). Moreover, Miller et al. (2019) note that first-order phase transitions are not allowed in the case of the spectral parameterization, but they are permitted non-exhaustively in the case of the PP expansion (as in the present work). Overall the priors on their parameterizations, in particular the spectral model (see also Figure 1) allow for a larger range of possible EOS functions. Visually comparing our inferences with Figure 14 in Miller et al. (2019), however, we deem the posterior distributions to be consistent.

Next, we compare to previous analyses of EOS constraints conditional on GW170817. The results of the LVC (Abbott et al. 2018) are broadly consistent with our analysis here, yet there are noticeable differences in the lower bound of the 90% confidence interval in radius (comparing Figures 3 and 4 with Figure 3 in Abbott et al. 2018). We attribute this discrepancy primarily to the different assumptions on the speed of sound: the spectral EOSs in Abbott et al. allowed models with up to 10% violation of causality, while both parameterized EOSs used in our analysis were strictly causal with $c_s \leq c$. Additional differences are that Abbott et al. use as the crust EOS a SLy model up to $\sim 0.5 \rho_s$ while we use the BPS model, and we incorporate nuclear physics constraints based on chiral EFT up

to around saturation density (Hebeler et al. 2013). The latter affects mainly the prior at the large radius end.

A number of independent analyses of the GW data also found broadly consistent results. De et al. (2018) analyzed the GW data with the source location and distance fixed to those determined from the electromagnetic observations, identified scaling relations between tidal parameters and mass ratio using PP models, and determined EOS and radius constraints for a $1.4 M_{\odot}$ star to $8.9 < R_{1.4} < 13.2$ km consistent with our results here. Most et al. (2018) computed a large catalog of PP EOSs, parameterizing both hadronic models and those with phase transitions. They analyzed the subset of these consistent with high-mass NSs and constraints on tidal deformability parameter $\tilde{\Lambda}$, both of which were imposed as hard cutoffs. Their results for the case $M_{\max} > 2.01 M_{\odot}$ and $\tilde{\Lambda}_{1.4} < 800$ (e.g., their Figure 1, top left panel) are consistent with our findings here. Capano et al. (2019) used a different speed of sound parameterization based on similar chiral EFT constraints, as well as different priors, in particular a uniform distribution in radius. Thus, their results using the GW data alone are skewed more toward smaller radii $9.2 \lesssim R_{1.4} \lesssim 12.3$ km when imposing the chiral EFT limits up to ρ_s . Essick et al. (2020) used a nonparametric EOS inference under different priors, also finding broadly consistent results.

In conclusion, the large statistical uncertainties in the available NICER and LIGO/Virgo likelihood functions lead to broad agreements on the EOS and radius constraints across different analyses. However, the impact of priors, assumptions, and parameterizations is starting to become discernible, as we have shown. Highly anticipated upcoming observations with LIGO/Virgo, NICER, and radio pulsars will constrain the EOS for populations of NSs, and yield unique insights into the properties of cold, dense matter. Our method can readily ingest the additional information from multiple sources, as well as incorporate new constraints from subatomic experiments and theory.

We thank the anonymous referee for comments that helped to improve this work. This work was supported in part by NASA through the NICER mission and the Astrophysics Explorers Program. G.R., T.H., and S.N. are grateful for support from the Nederlandse Organisatie voor Wetenschappelijk Onderzoek (NWO) through the VIDI and Projectruimte grants (PI Nissanke). T.E.R. and A.L.W. acknowledge support from ERC Starting grant No. 639217 CSINEUTRONSTAR (PI Watts). This work was sponsored by NWO Exact and Natural Sciences for the use of supercomputer facilities, and was carried out on the Dutch national e-infrastructure with the support of SURF Cooperative. S.K.G., K.H., and A.S. acknowledge support from the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)—Project-ID 279384907—SFB 1245. S.G. acknowledges the support of the CNES. R.M.L. acknowledges the support of NASA through Hubble Fellowship Program grant HST-HF2-51440.001. J.M.L. acknowledges support by NASA through the NICER mission with Grant 80NSSC17K0554 and by the U.S. DOE through Grant DE-FG02-87ER40317. This research has made extensive use of NASA’s Astrophysics Data System Bibliographic Services (ADS) and the arXiv. We thank Jocelyn Read, Wynn Ho, and Cole Miller for comments on a draft manuscript.

Software: Python/C language (Oliphant 2007), GNU Scientific Library (GSL; Gough 2009), NumPy (van der Walt et al. 2011),

Cython (Behnel et al. 2011), SciPy (Jones et al. 2001), MPI (Forum 1994), *MPI for Python* (Dalcín et al. 2008), Matplotlib (Hunter 2007; Droettboom et al. 2018), IPython (Perez & Granger 2007), Jupyter (Kluyver et al. 2016), MULTINEST (Feroz et al. 2009), PYMULTINEST (Buchner et al. 2014).

Appendix A GW Likelihood and Prior Implementation

The posterior distribution derived by the LVC, assuming both binary components are NSs, is

$$p(\mathbf{M}, \mathbf{\Lambda} | \mathbf{d}) \propto p(\mathbf{d} | \mathbf{M}, \mathbf{\Lambda}) \quad (\text{A1})$$

where $\mathbf{M} = [M_1, M_2]^\top$, $\mathbf{\Lambda} = [\Lambda_1, \Lambda_2]^\top$, \mathbf{d} is the strain data vector, and we omit the conditional argument representing the model. The joint prior density $p(\mathbf{M}, \mathbf{\Lambda})$ is flat on compact support, such that the posterior density is proportional to the nuisance-marginalized likelihood function $p(\mathbf{d} | \mathbf{M}, \mathbf{\Lambda})$. This function is symmetric under exchange of the binary components—i.e., $M_1 \leftrightarrow M_2$ and $\Lambda_1 \leftrightarrow \Lambda_2$. We are therefore free to define one mass as being associated with the most massive component: $M_i \geq M_j$ for $i \neq j$. If we opt *not* to, a numeric label is always associated with the same physical object. We are free to transform spaces, such that the likelihood function is $p(\mathbf{d} | q, M_{\text{chirp}}, \mathbf{\Lambda})$, where $q := M_2/M_1$, and M_{chirp} is a symmetric combination of M_1 and M_2 . Here, $q \in \mathbb{R}^+$ if no ordering of masses is enforced, in which case the nuisance-marginalized likelihood for $q = Q$ is equal to that for $q = 1/Q$ under exchange of properties. On the other hand, $0 < q \leq 1$ if M_1 is associated with the more massive component, which is the definition we choose.

The joint posterior distribution of interior parameters (omitting the global model as a conditional argument) is

$$p(\theta, \varepsilon | \mathbf{d}) \propto p(\mathbf{d} | \theta, \varepsilon) p(\varepsilon | \theta) p(\theta), \quad (\text{A2})$$

where θ are EOS parameters, and ε are central (energy) densities, such that $(\theta, \varepsilon_i) \mapsto (M_i, \Lambda_i)$. The EOS parameters also operate as hyperparameters, in the simplest mode as an upper bound on the prior support of $\varepsilon \in [a, b(\theta)]$. The population-level prior density of binaries $p(\varepsilon | \theta)$ is *assumed* to be separable when one density is associated with a particular physical component:

$$p(\varepsilon | \theta) = \prod_i p(\varepsilon_i | \theta), \quad (\text{A3})$$

where $p(\varepsilon_i | \theta)$ is identical $\forall i$. Let us assume $f(\varepsilon_i) \sim U(a, b)$ for $b = b(\theta)$, where $f(\varepsilon_i)$ is a monotone transformation such as a logarithm, which is our choice in the main text.

We decided to define one density parameter as that of the binary component with the highest central density (and thus total mass), and therefore $f(\varepsilon_1) \geq f(\varepsilon_2)$. Transforming the joint prior density function above, assuming $f(\varepsilon_i) = \varepsilon_i$ for notational simplicity, yields

$$p(\varepsilon_1, \varepsilon_2 | \theta) = p(\varepsilon_2 | \varepsilon_1, \theta) p(\varepsilon_1 | \theta) \quad (\text{A4})$$

where

$$p(\varepsilon_1 | \theta) = \frac{2(\varepsilon_1 - a)}{(b - a)^2} \quad \text{and} \quad p(\varepsilon_2 | \varepsilon_1) = \frac{1}{(\varepsilon_1 - a)}, \quad (\text{A5})$$

for support $\varepsilon_1 \in [a, b]$ and $\varepsilon_2 \in [a, \varepsilon_1]$. We then require (according to standard prior implementation for nested sampling)

$$x_1(\varepsilon_1; \theta) = \int_a^{\varepsilon_1} \frac{2(t - a)}{(b - a)^2} dt = \frac{(\varepsilon_1 - a)^2}{(b - a)^2}; \quad (\text{A6})$$

inverting yields

$$\varepsilon_1(x_1; \theta) = a + (b - a) \sqrt{x_1}. \quad (\text{A7})$$

Further,

$$\begin{aligned} x_2(\varepsilon_2; \varepsilon_1) &= \frac{1}{(\varepsilon_1 - a)} \int_a^{\varepsilon_2} dt = \frac{\varepsilon_2 - a}{(\varepsilon_1 - a)} \\ \Rightarrow \varepsilon_2(x_2; x_1) &= a + (\varepsilon_1 - a)x_2. \end{aligned} \quad (\text{A8})$$

Returning to the likelihood function $p(\mathbf{d} | \mathbf{M}, \mathbf{\Lambda})$, the function is almost degenerate, with support (defined in terms of some small threshold value of the likelihood normalized to the global maximum) only for $|g(\mathbf{M})| \leq \epsilon$, where $g(\mathbf{M})$ is the chirp combination minus some (now) constant M_{chirp} and $\epsilon \in \mathbb{R}^+$ is small. In the limit $\epsilon \rightarrow 0$, the likelihood function is degenerate in the \mathbf{M} -subspace. We enforce no order in M_1 and M_2 . Generally, the marginal posterior density is then

$$p(\theta, \varepsilon_1 | \mathbf{d}) \propto p(\varepsilon_1 | \theta) p(\theta) \times \int p(\mathbf{d} | \mathbf{M}, \mathbf{\Lambda}) p(\varepsilon_2 | \theta) d\varepsilon_2. \quad (\text{A9})$$

If $\epsilon \rightarrow 0$, then

$$\begin{aligned} p(\theta, \varepsilon_1 | \mathbf{d}) &\propto p(\varepsilon_1 | \theta) p(\theta) p(\mathbf{d} | M_1, \mathbf{\Lambda}) \\ &\times \int \delta(\varepsilon_2 - h(\varepsilon_1, \theta; M_{\text{chirp}})) p(\varepsilon_2 | \theta) d\varepsilon_2 \\ &= p(\varepsilon_1 | \theta) p(\theta) p(\mathbf{d} | M_1, \mathbf{\Lambda}) \underbrace{p(h(\varepsilon_1, \theta; M_{\text{chirp}}) | \theta)}_{\text{support } \varepsilon_2 \in [a, b]}, \end{aligned} \quad (\text{A10})$$

where M_2 is the total mass yielded via inversion of the chirp combination of \mathbf{M} , and $h(\varepsilon_1, \theta; M_{\text{chirp}})$ is the associated value of ε_1 for the EOS specified by θ . Further, $\Lambda_2 = \Lambda_2(\varepsilon_1, \theta; M_{\text{chirp}})$. The last factor implies that we need to take the prior support of ε_2 into account given hyperparameters θ ; note that if $h(\varepsilon_1, \theta; M_{\text{chirp}})$ has no solution, it will by definition not be within the prior support, and note also that the set of solutions to the inversion problem is generally a proper superset of the prior support.

However, because we enforced an order on the total masses, such that $M_1 \geq M_2$, then

$$p(\theta, \varepsilon_1 | \mathbf{d}) \propto p(\varepsilon_1 | \theta) p(\theta) \times \int p(\mathbf{d} | \mathbf{M}, \mathbf{\Lambda}) p(\varepsilon_2 | \varepsilon_1, \theta) d\varepsilon_2. \quad (\text{A11})$$

If $\epsilon \rightarrow 0$, then

$$\begin{aligned} p(\theta, \varepsilon_1 | \mathbf{d}) &\propto p(\varepsilon_1 | \theta) p(\theta) p(\mathbf{d} | M_1, \mathbf{\Lambda}) \\ &\times \int \delta(\varepsilon_2 - h(\varepsilon_1, \theta; M_{\text{chirp}})) p(\varepsilon_2 | \varepsilon_1, \theta) d\varepsilon_2 \\ &= p(\varepsilon_1 | \theta) p(\theta) p(\mathbf{d} | M_1, \mathbf{\Lambda}) \\ &\times \underbrace{p(h(\varepsilon_1, \theta; M_{\text{chirp}}) | \varepsilon_1, \theta)}_{\text{support } \varepsilon_2 \in [a, \varepsilon_1]}. \end{aligned} \quad (\text{A12})$$

This means that if the most massive binary component is such that $h(\varepsilon_1, \theta; M_{\text{chirp}}) \geq \varepsilon_1$, then the last factor imposes zero

(approximate) posterior density at point $[\theta, h(\varepsilon_1, \theta; M_{\text{chirp}})]^\top$. Note that if the hyperparameters θ merely delimit stability against radial perturbations, then $p(h(\varepsilon_1, \theta; M_{\text{chirp}})|\varepsilon_1, \theta)$ loses dependence, becoming $p(h(\varepsilon_1, \theta; M_{\text{chirp}})|\varepsilon_1)$. When transforming M_1 in Equation (A12) to the mass ratio q we obtain the GW-related part of Equation (3) in the main text.

In reality the likelihood function is finite for $|g(\mathbf{M})| \leq \epsilon$ with ϵ finite, with the marginal posterior distribution of the chirp combination having 90% credible interval about the median of $M_{\text{chirp}} = 1.186^{+0.001}_{-0.001} M_\odot$. If we are to account for this likelihood information accurately, but avoid sampling from a relatively very diffuse prior in the central densities and thus in the chirp mass, then we need to perform fast marginalization over the central density ε_2 to decrement the dimensionality of the sampling space. To proceed, we need to examine the *conditional* likelihood function

$$L(\varepsilon_2; \varepsilon_1, \theta) = p(\mathbf{d}|\mathbf{M}, \Lambda). \quad (\text{A13})$$

Given M_1 , there is one maximum in the conditional likelihood function when slicing through the chirp-sensitive likelihood function. Further, $M_2 = M_2(\varepsilon_2; \theta)$, while not a monotone transformation in the presence of unstable branches, is ordered, meaning there should exist one maximum with respect to ε_2 . In order to marginalize we therefore aim to first approximate the central density ε_2 that maximizes the conditional likelihood function. We thus approximate—via inversion— $h(\varepsilon_1, \theta; M'_{\text{chirp}})$, as above, where M'_{chirp} is now the median chirp mass in the marginal posterior distribution of M_{chirp} . Given this estimator of the maximum, we can generate bounds for numerical integration. For instance, we can integrate on the interval $\varepsilon_2 \in [\alpha, \beta]$, where $\alpha := h(\varepsilon_1, \theta; M'_{\text{chirp}} - nC_l)$ and $\beta := h(\varepsilon_1, \theta; M'_{\text{chirp}} + nC_u)$ where n is some integer and the X% marginal credible interval on the chirp mass is $\text{CI}_{X\%} = \{M_{\text{chirp}}; M_{\text{chirp}} \in [M'_{\text{chirp}} - C_l, M'_{\text{chirp}} + C_u]\}$. We can then perform numerical quadrature where the integrand requires numerical integration to transform $(\varepsilon_2, \theta) \mapsto (M_2, \Lambda_2)$. If an order is imposed on the central densities, then the upper-bound for quadrature is $\min[h(\varepsilon_1, \theta; M'_{\text{chirp}} + nC_u), \varepsilon_1]$.

Another approach would be to inject likelihood information into the *prior* density function of the mass M_2 (and thus of the central density ε_2). The prior *support* of M_2 (and thus ε_2) is restricted to some narrow interval corresponding to a narrow interval in chirp mass about the marginal median value. This is uncomfortable to have to rely on, one reason being that prior predictive probabilities may be compromised for model comparison; another is that a prior may be defined on the space of M_2 that is inconsistent with the prior density $p(\varepsilon_1|\theta)$, or a prior is defined directly on the space of the chirp mass, leading to a similar inconsistency. Given this modification of the prior support relative to the protocol outlined above, (nested) sampling proceeds without decrementing the dimensionality, but at far lower cost and effectively without risk of error due to insufficient resolution.

In this work we compared two approaches quantitatively: (i) the delta-function approximation—specifically Equation (A12), and (ii) the approach wherein likelihood information is injected into the prior via modification of the support, in the form of a narrow prior on the chirp mass about the median in marginal

posterior mass; see the text following Equation (3). The joint posterior distribution of the EOS parameters is consistent for both approaches.

Appendix B Approximations to the Likelihood Function

B.1. The NS–NS Scenario

The posterior distribution on the two binary components from GW170817 derived by the LVC (see Abbott et al. 2019a) is related to the nuisance-marginalized likelihood function through

$$p(\Lambda_1, \Lambda_2, M_1, M_2 | \mathbf{d}_{\text{GW}}) \propto p(\Lambda_1, \Lambda_2, M_1, M_2) p(\mathbf{d}_{\text{GW}} | \Lambda_1, \Lambda_2, M_1, M_2), \quad (\text{B1})$$

where we have only shown the four parameters of interest here. In this work we have equated the two by arguing that the prior distribution $p(\Lambda_1, \Lambda_2, M_1, M_2)$ is jointly flat. Furthermore we have performed a coordinate transformation from the two component masses to the mass ratio q and the chirp mass M_{chirp} to get

$$p(\Lambda_1, \Lambda_2, q, M_{\text{chirp}} | \mathbf{d}_{\text{GW}}). \quad (\text{B2})$$

As the chirp mass is constrained to a very small range we fix its value to $M_{\text{chirp}} = 1.186 M_\odot$ such that we can write a conditional distribution

$$p(\Lambda_1, \Lambda_2, q | \mathbf{d}_{\text{GW}}, M_{\text{chirp}}), \quad (\text{B3})$$

which is a slice through the posterior distribution. We use, however, the marginalized posterior distribution $p(\Lambda_1, \Lambda_2, q | \mathbf{d}_{\text{GW}})$ in this Letter, which we compare to the conditional distribution in the left panel of Figure B1.

We furthermore note that when transforming from component masses to mass ratio and chirp mass the prior on these quantities is no longer jointly flat. We therefore check that the posterior distribution on q does not change significantly when applying a weighting to the distribution to ensure a jointly flat prior. We show the difference between the distributions in the right panel of Figure B1.

B.2. The NS–BH Scenario

Finally we have approximated the likelihood function for the NS–BH case by transforming Λ_1 and Λ_2 to the parameter $\tilde{\Lambda}$ and marginalizing over Λ_1 , which is different from taking a slice through the posterior at $\Lambda_1 = 0$. However, to compute the constraints from the EM counterpart the latter is not possible since there are not enough samples in the posterior that have $\Lambda_1 = 0$. As an alternative we approximate $\Lambda_1 = 0$ by taking the samples in the distribution where $\Lambda_1 < 30$, such that we are not dominated by stochastic noise of having too few samples. In Figure B2 we compare the two approaches when we set the posterior support for parameter sets that result in $M_{\text{rem}} < 0.1 M_\odot$ to zero, while setting uniform support everywhere else. Including these constraints from the EM counterpart, we conclude that there is a negligible change in the posterior distribution of $\tilde{\Lambda}$.




Figure B1. Left panel: comparison of the distribution of the mass ratio q when we use the posterior distribution marginalized over the chirp mass (green) or the posterior distribution conditional on a given chirp mass (blue). Both distributions are very similar, which is why we argue that the usage of the marginalized distribution throughout the Letter is justified. Right panel: prior distribution on the mass ratio q in the black dashed lines, the posterior distribution on q in green, and the reweighted posterior that corresponds to a flat prior in q in blue. Note that the prior distribution on q is only for a small interval in chirp mass where there is posterior support, which is why there is more support toward lower- than higher-mass ratios. We conclude that there is slightly more posterior support for equal mass ratios when reweighting but this is negligible compared to the statistical uncertainty in GW170817.




Figure B2. Comparison between the posterior distribution on $\tilde{\Lambda}$ when marginalizing over Λ_1 and when only considering values of $\Lambda_1 < 30$. The range $\Lambda_1 < 30$ is chosen in order to well-approximate the conditional distribution at $\Lambda_1 = 0$ while reducing noise in the histogram. The distributions shown here have already implemented zero support on parameter sets that result in $M_{\text{rem}} < 0.1 M_{\odot}$ and uniform support where $M_{\text{rem}} > 0.1$ based on the electromagnetic analysis in Section 3.2 and Hinderer et al. (2019). The difference between the two distributions is small enough to justify the approximation made in this Letter.

ORCID iDs

G. Raaijmakers <https://orcid.org/0000-0002-9397-786X>
T. E. Riley <https://orcid.org/0000-0001-9313-0493>
T. Hinderer <https://orcid.org/0000-0002-3394-6105>
A. L. Watts <https://orcid.org/0000-0002-1009-2354>
S. Nissanke <https://orcid.org/0000-0001-6573-7773>
S. Guillot <https://orcid.org/0000-0002-6449-106X>
R. M. Ludlam <https://orcid.org/0000-0002-8961-939X>

References

Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017, *PhRvL*, **119**, 161101
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2018, *PhRvL*, **121**, 161101

Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019a, *PhRvX*, **9**, 011001
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019b, *PhRvX*, **9**, 031040
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2020, *ApJL*, **892**, L3
Alsing, J., Silva, H. O., & Berti, E. 2018, *MNRAS*, **478**, 1377
Alvarez-Castillo, D., Ayriyan, A., Benic, S., et al. 2016, *EPJA*, **52**, 69
Annala, E., Gorda, T., Kurkela, A., & Vuorinen, A. 2018, *PhRvL*, **120**, 172703
Antoniadis, J., Freire, P. C. C., Wex, N., et al. 2013, *Sci*, **340**, 448
Arzoumanian, Z., Brazier, A., Burke-Spolaor, S., et al. 2018, *ApJS*, **235**, 37
Ascenzi, S., Lillo, N. D., Haster, C.-J., Ohme, F., & Pannarale, F. 2019, *ApJ*, **877**, 94
Bauswein, A., Just, O., Janka, H.-T., & Stergioulas, N. 2017, *ApJL*, **850**, L34
Baym, G., Hatsuda, T., Kojo, T., et al. 2018, *RPPh*, **81**, 056902
Baym, G., Pethick, C., & Sutherland, P. 1971, *ApJ*, **170**, 299
Behnel, S., Bradshaw, R., Citro, C., et al. 2011, *CSE*, **13**, 31
Buchner, J., Georgakakis, A., Nandra, K., et al. 2014, *A&A*, **564**, A125
Capano, C. D., Tews, I., Brown, S. M., et al. 2019, *NatAs*, in press (arXiv:1908.10352)
Coughlin, M. W., & Dietrich, T. 2019, *PhRvD*, **100**, 043011
Coughlin, M. W., Dietrich, T., Doctor, Z., et al. 2018, *MNRAS*, **480**, 3871
Cromartie, H. T., Fonseca, E., Ransom, S. M., et al. 2019, *NatAs*, **4**, 72
Dalcín, L., Paz, R., Storti, M., & D'Elia, J. 2008, *JPCD*, **68**, 655
De, S., Finstad, D., Lattimer, J. M., et al. 2018, *PhRvL*, **121**, 091102
Demorest, P. B., Pennucci, T., Ransom, S. M., Roberts, M. S. E., & Hessels, J. W. T. 2010, *Natur*, **467**, 1081
Douchin, F., & Haensel, P. 2001, *A&A*, **380**, 151
Droettboom, M., Caswell, T. A., Hunter, J., et al. 2018, matplotlib/matplotlib v2.2.2, Zenodo, doi:10.5281/zenodo.1202077
Essick, R., Landry, P., & Holz, D. E. 2020, *PhRvD*, **101**, 063007
Feroz, F., & Hobson, M. P. 2008, *MNRAS*, **384**, 449
Feroz, F., Hobson, M. P., & Bridges, M. 2009, *MNRAS*, **398**, 1601
Feroz, F., Hobson, M. P., Cameron, E., & Pettitt, A. N. 2013, arXiv:1306.2144
Fonseca, E., Pennucci, T. T., Ellis, J. A., et al. 2016, *ApJ*, **832**, 167
Forum, M. P. 1994, MPI: A Message-Passing Interface Standard (New York: Association for Computing Machinery)
Foucart, F., Hinderer, T., & Nissanke, S. 2018, *PhRvD*, **98**, 081501
Fraga, E. S., Kurkela, A., & Vuorinen, A. 2014, *ApJL*, **781**, L25
Gao, H., Cao, Z., Ai, S., & Zhang, B. 2017, *ApJL*, **851**, L45
García-Bellido, J., Linde, A., & Wands, D. 1996, *PhRvD*, **54**, 6040
Gough, B. 2009, GNU Scientific Library Reference Manual (3rd ed.; Boston, MA: Network Theory Ltd.)
Greif, S. K., Raaijmakers, G., Hebeler, K., Schwenk, A., & Watts, A. L. 2019, *MNRAS*, **485**, 5363
Hebeler, K., Holt, J. D., Menéndez, J., & Schwenk, A. 2015, *ARNPS*, **65**, 457
Hebeler, K., Lattimer, J. M., Pethick, C. J., & Schwenk, A. 2013, *ApJ*, **773**, 11
Hebeler, K., & Schwenk, A. 2010, *PhRvC*, **82**, 014314
Hinderer, T., Nissanke, S., Foucart, F., et al. 2019, *PhRvD*, **100**, 063021
Hunter, J. D. 2007, *CSE*, **9**, 90

- Jones, E., Oliphant, T., Peterson, P., et al. 2001, SciPy: Open Source Scientific Tools for Python, Online; accessed 21.06.2019, <http://www.scipy.org/>
- Kasliwal, M. M., Nakar, E., Singer, L. P., et al. 2017, *Sci*, **358**, 1559
- Kass, R. E., & Raftery, A. E. 1995, *Journal of the American Statistical Association*, **90**, 773
- Kastaun, W., & Ohme, F. 2019, *PhRvD*, **100**, 103023
- Kluyver, T., Ragan-Kelley, B., Pérez, F., et al. 2016, in *Positioning and Power in Academic Publishing: Players, Agents and Agendas*, ed. F. Loizides & B. Schmidt (Amsterdam: IOS Press), 87
- Kulkarni, S. R. 2005, arXiv:[astro-ph/0510256](https://arxiv.org/abs/astro-ph/0510256)
- Kullback, S., & Leibler, R. A. 1951, *Ann. Math. Statist.*, **22**, 79
- Lattimer, J. M., & Prakash, M. 2016, *PhR*, **621**, 127
- Lattimer, J. M., & Schramm, D. N. 1976, *ApJ*, **210**, 549
- Li, L.-X., & Paczyński, B. 1998, *ApJL*, **507**, L59
- Lindblom, L. 2010, *PhRvD*, **82**, 103011
- Lindblom, L. 2018, *PhRvD*, **97**, 123019
- Margalit, B., & Metzger, B. D. 2017, *ApJL*, **850**, L19
- Margalit, B., & Metzger, B. D. 2019, *ApJL*, **880**, L15
- Metzger, B. D. 2017, *LRR*, **20**, 3
- Metzger, B. D., Martínez-Pinedo, G., Darbha, S., et al. 2010, *MNRAS*, **406**, 2650
- Miller, M. C., Chirenti, C., & Lamb, F. K. 2020, *ApJ*, **888**, 12
- Miller, M. C., Lamb, F. K., Dittmann, A. J., et al. 2019, *ApJL*, **887**, L24
- Most, E. R., Weih, L. R., Rezzolla, L., & Schaffner-Bielich, J. 2018, *PhRvL*, **120**, 261103
- Oertel, M., Hempel, M., Klähn, T., & Typel, S. 2017, *RvMP*, **89**, 015007
- Oliphant, T. E. 2007, *CSE*, **9**, 10
- Perez, F., & Granger, B. E. 2007, *CSE*, **9**, 21
- Raaijmakers, G., Greif, S. K., Riley, T. E., et al. 2020, Constraining the Dense Matter Equation of State with Joint Analysis of NICER and LIGO/Virgo Measurements: Data for Generating Plots, Zenodo, doi:[10.5281/zenodo.3711718](https://doi.org/10.5281/zenodo.3711718)
- Raaijmakers, G., Riley, T. E., Watts, A. L., et al. 2019, *ApJL*, **887**, L22
- Radice, D., & Dai, L. 2019, *EPJA*, **A55**, 50
- Rezzolla, L., Most, E. R., & Weih, L. R. 2018, *ApJL*, **852**, L25
- Riley, T. E., Raaijmakers, G., & Watts, A. L. 2018, *MNRAS*, **478**, 1093
- Riley, T. E., Watts, A. L., Bogdanov, S., et al. 2019a, *ApJL*, **887**, L21
- Riley, T. E., Watts, A. L., Bogdanov, S., et al. 2019b, A NICER View of PSR J0030+0451: Nested Samples for Millisecond Pulsar Parameter Estimation, v1.1.0, Zenodo, doi:[10.5281/zenodo.3707821](https://doi.org/10.5281/zenodo.3707821)
- Rosswog, S., Liebendörfer, M., Thielemann, F.-K., et al. 1999, *A&A*, **341**, 499
- Ruiz, M., Shapiro, S. L., & Tsokaros, A. 2018, *PhRvD*, **97**, 021501
- Shibata, M., Fujibayashi, S., Hotokezaka, K., et al. 2017, *PhRvD*, **96**, 123012
- Shibata, M., Zhou, E., Kiuchi, K., & Fujibayashi, S. 2019, *PhRvD*, **100**, 023015
- Tews, I., Margueron, J., & Reddy, S. 2018, *PhRvC*, **98**, 045804
- Tews, I., & Schwenk, A. 2020, *ApJ*, **892**, 14
- The LIGO Scientific Collaboration & The Virgo Collaboration 2018, GWTC-1 Public Data Release, <https://dcc.ligo.org/LIGO-P1800370/public>
- van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, *CSE*, **13**, 22
- Watts, A. L. 2019, in *AIP Conf. Ser. 2127, Xiamen-CUSTIPEN Workshop on the Equation of State of Dense Neutron-Rich Matter in the Era of Gravitational Wave Astronomy*, ed. A. Li, B.-A. Li, & F. Fu (Melville, NY: AIP), 020008
- Yang, H., East, W. E., & Lehner, L. 2018, *ApJ*, **856**, 110