

This is the accepted manuscript made available via CHORUS. The article has been published as:

Nonparametric inference of the neutron star equation of state from gravitational wave observations

Philippe Landry and Reed Essick

Phys. Rev. D **99**, 084049 — Published 30 April 2019

DOI: [10.1103/PhysRevD.99.084049](https://doi.org/10.1103/PhysRevD.99.084049)

Non-parametric inference of the neutron star equation of state from gravitational wave observations

Philippe Landry*

*Enrico Fermi Institute and Kavli Institute for Cosmological Physics,
The University of Chicago, 5640 South Ellis Avenue, Chicago, Illinois, 60637, USA*

Reed Essick†

*Kavli Institute for Cosmological Physics, The University of Chicago,
5640 South Ellis Avenue, Chicago, Illinois, 60637, USA*

(Dated: February 20, 2019)

We develop a non-parametric method for inferring the universal neutron star (NS) equation of state (EOS) from gravitational wave (GW) observations. Many different possible realizations of the EOS are generated with a Gaussian process conditioned on a set of nuclear-theoretic models. These synthetic EOSs are causal and thermodynamically stable by construction, span a broad region of the pressure-density plane, and can be selected to satisfy astrophysical constraints on the NS mass. Associating every synthetic EOS with a pair of component masses $M_{1,2}$ and calculating the corresponding tidal deformabilities $\Lambda_{1,2}$, we perform Monte Carlo integration over the GW likelihood for $M_{1,2}$ and $\Lambda_{1,2}$ to directly infer a posterior process for the NS EOS. We first demonstrate that the method can accurately recover the properties of an injected GW signal, and subsequently use it to analyze data from GW170817, finding a canonical deformability of $\Lambda_{1.4} = 160^{+48}_{-113}$ and $p(2\rho_{\text{nuc}}) = 1.35^{+1.8}_{-1.2} \times 10^{34}$ dyn/cm² for the pressure at twice the nuclear saturation density at 90% confidence, in agreement with previous studies, when assuming a loose EOS prior. With a prior more tightly constrained to resemble the theoretical EOS models, we recover $\Lambda_{1.4} = 556^{+163}_{-172}$ and $p(2\rho_{\text{nuc}}) = 4.73^{+1.4}_{-2.5} \times 10^{34}$ dyn/cm². We further infer the maximum NS mass supported by the EOS to be $M_{\text{max}} = 2.09^{+0.37}_{-0.16}$ ($2.04^{+0.22}_{-0.002}$) M_{\odot} with the loose (tight) prior. The Bayes factor between the two priors is $B_{\mathcal{T}}^A \simeq 1.12$, suggesting that neither is strongly preferred by the data and that constraints on the EOS from GW170817 alone may be sensitive to the choice of prior.

I. INTRODUCTION

Determining the neutron star (NS) equation of state (EOS) is a major unsolved problem in nuclear astrophysics. The NS EOS is the zero-temperature pressure-density relation that arises from the star's nuclear microphysics, encoding its composition and internal structure, and determining its macroscopic properties. Because young NSs cool rapidly via neutrino emission and quickly reach β -equilibrium, thermal and dissipative corrections to the EOS are negligible, and NS matter is typically modeled as a perfect fluid [1]. The EOS for the low-density NS crust is well known [2], but the conditions in the dense NS core are so extreme that laboratory experiments are unable to probe its constituent supranuclear matter. Consequently, astrophysical observations of NSs offer the best opportunity for constraining the unknown core EOS.

Because a given candidate EOS prescribes a unique mass-radius relation [3], simultaneous mass and radius measurements have been attempted for a number of NSs [4]. However, obtaining an accurate radius measurement, which involves modeling the star's thermal X-ray emission, is notoriously challenging [5]. Likewise, simultaneous pulsar mass and moment of inertia measurements have been proposed as a means of constraining the NS EOS via the mass-moment of inertia relation [6, 7]. Existing radio observatories are, however, unable to measure the relativistic periastron advance of any known binary pulsar with sufficient precision to infer the stellar moment of inertia [8]. Although future moment of inertia measurements from next-generation radio telescopes, such as the Square Kilometre Array [9], and forthcoming radius measurements from the NICER soft X-ray observatory [10] are expected to make significant contributions to the study of ultra-dense matter, at present the most informative constraints on the NS EOS come from gravitational-wave (GW) astronomy. Indeed, Advanced LIGO's [11] and Virgo's [12] GW measurement of NS tidal deformability in GW170817 [13], the loud signal of a binary NS merger, was the basis for the analyses that produced the most discriminating constraints on the supranuclear EOS to date [14, 15].

The tidal forces that arise in NS binaries during inspiral deform the stars away from their spherical equilibrium shape. The mass quadrupole moments they acquire draw energy from the orbit and add to the binary's gravitational

* landryp@uchicago.edu

† reedessick@kicp.uchicago.edu

radiation, enhancing its GW luminosity and leading to a slight acceleration of the coalescence. This manifests as a phase shift in the waveform relative to the merger of point-particles [16–18], which is measurable with existing GW detectors [19–23]. Its magnitude depends on the size of the induced stellar quadrupoles, measured by the tidal deformabilities $\Lambda_{1,2}$ of the two NSs. Λ is an EOS-dependent, dimensionless function of the NS mass M that correlates with the pressure gradients inside the star—i.e. with the stiffness of the EOS. A stiff EOS has large pressure gradients that support bigger, more diffuse stars with large Λ , while a soft EOS has small pressure gradients that yield more compact stars with small Λ . Thus, as with the mass-radius or mass-moment of inertia relation, simultaneous measurements can pick out the preferred M - Λ relation, which is in one-to-one correspondence with the NS EOS.

The first GW constraints on NS tidal deformability were reported in Ref. [13]. The parameters $\Lambda_{1,2}$ were included in the waveform model and inferred, along with the masses $M_{1,2}$ and the other source properties, from GW170817 strain data [24]. The resulting posterior distribution over the tidal deformabilities and masses established an upper bound on the chirp deformability

$$\tilde{\Lambda} = \frac{16}{13} \frac{(M_1 + 12M_2)M_1^4\Lambda_1 + (M_2 + 12M_1)M_2^4\Lambda_2}{(M_1 + M_2)^5}, \quad (1)$$

the particular mass-weighted average of $\Lambda_{1,2}$ that appears in the waveform at lowest post-Newtonian order [18, 25]. Assuming that GW170817’s components rotated slowly [23, 26, 27], with dimensionless spin $\chi_{1,2} := cS_{1,2}/GM_{1,2}^2 \leq 0.05$, Ref. [13] placed a 90%-credible upper bound $\tilde{\Lambda} \leq 800$. The upper limit on the deformability of a canonical $1.4 M_\odot$ NS was also found to be $\Lambda_{1.4} := \Lambda(1.4 M_\odot) \leq 800$. This was refined to $\tilde{\Lambda} = 300^{+420}_{-230}$ (median and highest-posterior-density 90% confidence interval) in a subsequent analysis [28], which made use of improved waveform models and included lower frequencies within the GW data. For comparison, the widely used candidate EOS **slly** [29], among the softer theoretical models, has $\Lambda_{1.4} \approx 290$, while **ms1b** [30], one of the stiffest, has $\Lambda_{1.4} \approx 1220$. Overall, these constraints favor a relatively soft NS EOS, and bounds from Ref. [28] disfavor several candidate EOSs at the 90% confidence level.

Nonetheless, the analyses of Refs. [13, 28] neglected important correlations between Λ_1 and Λ_2 and did not attempt to translate the tidal constraints into a direct inference of the NS EOS. Because a NS’s composition is dictated by universal nuclear many-body physics, every NS is thought to share a common EOS. Their individual properties can thus be parameterized by the stellar mass alone. In particular, since all NSs conform to the same M - Λ relation, the waveform parameters $\Lambda_{1,2}$ are not truly independent. For example, with the standard convention $M_2 \leq M_1$, $\Lambda_2 \geq \Lambda_1$ for any physically realistic EOS, as $\Lambda(M)$ decreases monotonically. Moreover, the correlation between $\Lambda_{1,2}$ has been shown to be approximately universal [31, 32]; knowledge of Λ_1 and the binary mass ratio is sufficient to determine Λ_2 with high fidelity, regardless of the EOS. Deviations from the universal relation are at most $\sim 10\%$, and are substantially smaller for nearly equal-mass systems. Incorporating these correlations, known as *binary Love relations*, into a prior on $\Lambda_{1,2}$ can significantly improve the recovery of tidal information from GWs. An injection study performed in Ref. [33] demonstrated that the area of the 90%-credible region of the Λ_1 - Λ_2 posterior is reduced by a factor of two or more.

The intrinsic correlations captured by the binary Love relations were taken into account in a pair of reanalyses of GW170817. Ref. [15] made use of the EOS-insensitive, mass-ratio dependent binary Love relations from Refs. [31, 32] to compute Λ_2 as a function of the other parameters in the waveform model. Ref. [14] used a closely related approximation, $\Lambda_2 = (M_1/M_2)^6\Lambda_1$, for the same purpose. With the inclusion of these physical restrictions on the tidal deformabilities, both works tightened the tidal constraints inferred from GW170817. Ref. [14] reported $\tilde{\Lambda} = 222^{+420}_{-138}$ for a uniform component-mass prior, while Ref. [15] found $\Lambda_{1.4} = 190^{+390}_{-120}$ (both median and symmetric 90% confidence interval). Both studies employed the low-spin prior described above, and they prefer a marginally softer EOS than was originally inferred.

Ref. [15] also sought to constrain the NS EOS directly, instead of working exclusively with the tidal deformabilities. It adopted a spectral parameterization for the EOS [34–37] following methodology originally developed for piecewise polytropes [38–41]. The spectral coefficients were sampled in place of $\Lambda_{1,2}$ in the waveform, and the most probable EOS was then reconstructed, with error bars, from the first four spectral coefficients. This *spectral method* ultimately produced a *process* in the pressure-density plane, in addition to a joint posterior on $\Lambda_{1,2}$; just as a distribution gives the relative probability of finite sets of variates, a process gives the relative probability of functional degrees of freedom. Although the $\Lambda_{1,2}$ posterior can be used for Bayesian hypothesis ranking of different candidate EOSs [21, 42, 43], the EOS process is inherently more informative. The discrete set of nuclear theory models tested in a hypothesis ranking scheme may not capture the full range of possibilities for the EOS, and reducing each model to an overall evidence score may make it difficult to deduce preferences for specific fine-grained features of the EOS. In contrast, Ref. [15] was able to report, e.g., the recovered pressure at twice and six times the nuclear saturation density ($\rho_{\text{nuc}} = 2.8 \times 10^{14} \text{ g/cm}^3$) independently of existing nuclear-theoretic models with $\sim 60\%$ uncertainty.

However, parametric EOS inference has limitations of its own. The reduction of a complicated pressure-density

function to a small number of parameters necessarily leads to modeling errors [41]. This problem is especially acute for sharp features in the pressure-density relation, such as first-order phase transitions, and a four-parameter spectral representation has been shown to model such discontinuous behavior rather poorly [35]. Indeed, it performs no better in this regard than a three-segment piecewise polytrope parameterization, whose issues with phase transitions are widely acknowledged [35, 39]. While it may be possible to design an alternate parameterization specially adapted to such phase transitions, given the extremely varied phenomenology of hybrid EOSs (e.g. [44–47]; see Ref. [48] for an overview), it is unlikely that any single parametric model will be able to faithfully represent the full range of EOS variability with only a handful of parameters. Accordingly, a *non-parametric* representation of the EOS may be better suited to the task of inferring the internal structure of NSs from astrophysical observations.

The faithfulness of a non-parametric representation of a function f scales with the amount of underlying knowledge of f —the way, e.g., the appropriate number of bins to use in a histogram grows as more data are observed—instead of the number of parameters chosen in a parametric model. Non-parametric representations typically involve a set of *hyperparameters* that control allowed types of functional behavior—like the number and placement of bins in a histogram—but the resulting freedom in f is much larger than that afforded by a parametric model. This is because hyperparameters describe the correlations between a function’s values rather than the values themselves. Hence, the uncertainty in the non-parametric representation decreases as observations of the function accrue, while systematic errors in the parametric representation are fixed by the degree to which the model is suited to f . Given that information about the unknown EOS will accumulate through successive GW observations, a non-parametric approach is well-adapted to the so-called relativistic inverse stellar structure problem.

In this paper, we introduce a non-parametric method for directly inferring the NS EOS from GW data. Rather than prescribing a functional form for the relation between the star’s pressure p and its total energy density ε , as was done in Ref. [15], we use Gaussian process regression (GPR) to generate a large number of possible realizations $\varepsilon^{(a)}(p)$ of the EOS, labeled by $a = 1, 2, \dots, N$. These *synthetic EOSs* span the full range of stiffnesses, core pressures and other characteristics consistent with thermodynamic stability, causality, astrophysical observations of NSs, and candidate models $\varepsilon^{(\alpha)}(p)$, $\alpha = 1, 2, \dots, n$, from nuclear theory.¹ We then associate each synthetic EOS $\varepsilon^{(a)}(p)$ with NS masses $M_{1,2}^{(a)}$ drawn from a prior distribution and calculate the corresponding tidal deformabilities $\Lambda_{1,2}^{(a)}$. A Gaussian kernel density estimate (KDE) approximates the marginal conditional likelihood $\mathcal{L}^{(a)} := \mathcal{L}(d|M_1^{(a)}, M_2^{(a)}, \Lambda_1^{(a)}, \Lambda_2^{(a)}; \mathcal{H})$ obtained from standard parameter estimation [24, 49]. By assigning $\mathcal{L}^{(a)}$ to $\varepsilon^{(a)}(p)$, we establish a statistical map between observables $(M_{1,2}, \Lambda_{1,2})$ and the pressure-density plane. In the limit of many samples ($N \rightarrow \infty$), we obtain a smooth *EOS posterior process*. This inference scheme is similar to what was used in Ref. [39] in the context of piecewise polytropes, but our approach adopts a different representation of the EOS prior and employs different techniques for obtaining the posterior. Ref. [50] employs integration techniques similar to ours in order to bound macroscopic NS observables, but adopts an *ad hoc* EOS prior and applies cuts based on the $\Lambda_{1.4}$ bounds reported in Refs. [13, 51].

Besides possessing systematics that are completely independent from those for parametric methods, a non-parametric approach to EOS inference has several attractive features. Instead of requiring a separate analysis of GW strain data with an alternate waveform model, our method operates on the standard GW $M_{1,2}$ - $\Lambda_{1,2}$ likelihoods produced by parameter estimation. Moreover, a GPR representation of the EOS naturally comes equipped with smooth error estimates for the function, unlike simple interpolation or a spectral construction. Most importantly, such a representation provides immense flexibility when selecting the EOS prior; by conditioning the GPR on different sets of candidate EOSs, priors specialized to different classes of theoretical models can easily be specified. With these, one could investigate, e.g., the prevalence of hyperonic degrees of freedom or first-order phase transitions in the EOS by calculating the relative posterior support for the relevant priors. Astrophysical priors on NS masses and radii can also be accommodated by the GPR, and physical constraints on the EOS—like causality, thermodynamic stability, and the binary Love relations—are automatically incorporated. Furthermore, conditioning the priors on tabulated EOSs may make our prior beliefs more transparent compared to *ad hoc* choices for parameters which may not have immediate physical interpretations. The flexibility and transparency afforded by a non-parametric representation of the EOS may become increasingly important as more binary NS systems are observed, reducing statistical uncertainty to levels comparable to systematic model uncertainties [52].

Leveraging these advantages, we develop two non-parametric priors with Gaussian processes conditioned on a fiducial set of tabulated candidate EOSs. Our *model-agnostic prior* attempts to span the full range of physically plausible EOSs and is only loosely informed by the theoretical models, while our *model-informed prior* is more tightly constrained by the candidate EOSs. Applying both priors to a series of simulated signals, we confirm that our inference scheme behaves as expected. Performing our analysis on GW170817, we infer $\bar{\Lambda} = 210_{-113}^{+383}$

¹ We index tabulated candidate EOSs with Greek letters $(\alpha), (\beta), (\gamma), \dots$ and synthetic EOSs with Latin letters $(a), (b), (c), \dots$.

(631_{-122}^{+164}) and $\Lambda_{1.4} = 160_{-113}^{+448}$ (556_{-172}^{+163}) with the model-agnostic (model-informed) prior.² We also infer the pressures $p(2\rho_{\text{nuc}}) = 1.35_{-1.2}^{+1.8} \times 10^{34}$ ($4.73_{-2.5}^{+1.4} \times 10^{34}$) dyn/cm² and $p(6\rho_{\text{nuc}}) = 8.86_{-5.9}^{+4.3} \times 10^{35}$ ($7.55_{-3.2}^{+2.0} \times 10^{35}$) dyn/cm². Both the model-agnostic and model-informed results are consistent with previous analyses of GW170817 [14, 15]. Additionally, we extrapolate the correlations between mid-range and high densities in the synthetic EOSs to place the first constraints on the maximum NS mass derived from GW data alone. Our analysis with the model-agnostic (model-informed) prior yields $M_{\text{max}} = 2.09_{-0.16}^{+0.37}$ ($2.04_{-0.002}^{+0.22}$) M_{\odot} . Computing the Bayes factor $B_{\mathcal{T}}^{\mathcal{A}}$ between the two priors, we find that neither is strongly preferred by the data: $B_{\mathcal{T}}^{\mathcal{A}} = 1.12 \pm 0.06$ (point estimate and 1- σ uncertainty). The fact that the inferred constraints vary significantly depending on the choice of prior, while the Bayes factor is near unity, suggests that constraints on the EOS and tidal deformabilities from GW170817 may be influenced by the choice of prior. Modeling systematics associated with the parameterization of prior beliefs may therefore be important.

We detail our development of Gaussian processes in Sec. II and describe our inference scheme in Sec. III. Sec. IV reports the specific priors used in this analysis. We apply them to simulated signals in Sec. V and to real GW170817 data in Sec. VI. We present our conclusions in Sec. VII.

II. NON-PARAMETRIC REPRESENTATION OF THE EQUATION OF STATE

Our non-parametric approach to GW-based EOS inference relies heavily on Gaussian processes (GPs). We therefore begin by reviewing GPs in general, before explaining our specific implementation for causal, thermodynamically stable EOSs. GPs have been applied to many problems in GW astronomy, such as the modeling of gravitational waveforms [53, 54] and electromagnetic counterpart light curves [55], the optimization of parameter estimation strategies [56–58], and hierarchical population inference [59, 60], but they have not previously been used in the relativistic inverse stellar structure problem. Furthermore, many of these applications only use GPs to interpolate between a sparse sample of known functions, marginalizing over interpolation uncertainty, rather than employing a GP directly as a Bayesian prior. We use GPs for both purposes. A helpful general reference about GPs can be found in Ref. [61].

A. Gaussian processes

GPs provide an extremely flexible and compact representation of the uncertainty in a function’s values (ordinates). A GP treats a real-valued function f as an element of an infinite-dimensional vector space with correlations between its ordinates $f_i := f(x_i) \in (-\infty, \infty)$ modeled as a multivariate Gaussian distribution \mathcal{N} with mean $\langle f_i \rangle$ and covariance $\text{Cov}(f_i, f_j)$. The joint distribution on the function’s ordinates, conditioned on their corresponding abscissae x_i , is thus

$$f_i | x_i \sim \mathcal{N}(\langle f_i \rangle, \text{Cov}(f_i, f_j)). \quad (2)$$

The elements of the covariance matrix are determined by a covariance kernel K such that $K_{ij} := K(x_i, x_j) = \text{Cov}(f_i, f_j)$. Typically, K is chosen such that the covariance is larger for abscissae that are closer together, effectively constraining f to be smooth over a length scale set by the kernel. More generally, K determines the preferred functional behavior of realizations drawn from the GP. Each realization f of the GP is therefore a sequence of correlated random variables, but all such functions nonetheless share some general features.

A standard choice for K is the *squared-exponential kernel*

$$K_{\text{se}}(x, x') = \sigma^2 \exp\left(-\frac{(x - x')^2}{2l^2}\right), \quad (3)$$

in which the hyperparameters σ and l determine the functions’ behavior: σ sets the overall strength of the correlations, and l governs the length scale over which they occur. Although we did not investigate alternative choices, other common kernels are the Ornstein-Uhlenbeck and Mortén kernels (see Sec. 4.2 of Ref. [61] for a more complete list). We choose the squared-exponential because it is analytic and adequately captures the observed variations within our data. Hence, our GP is built from this kernel, along with variants of the white-noise kernel

$$K_{\text{wn}}(x, x') = \sigma_{\text{wn}}^2(x) \delta(x - x'), \quad (4)$$

² Here and in the remainder of the paper, quoted error bars refer to highest-posterior-density 90% confidence intervals about the maximum *a posteriori* unless otherwise specified.

which models measurement uncertainty at a specific point.

One can determine conditional distributions for an arbitrary set of ordinates f_i given a set $\{f_{j^*}\}$ of known values via

$$P(f_i | x_i, \{f_{j^*}, x_{j^*}\}) = \frac{P(f_i, \{f_{j^*}\} | x_i, \{x_{j^*}\})}{P(\{f_{j^*}\} | \{x_{j^*}\})}, \quad (5)$$

the definition of a conditional probability. Here and throughout this work, starred indices (e.g. x_{i^*}) indicate abscissae upon which we condition the GP, as opposed to generic abscissae (e.g. x_i) at which the ordinates are unknown. We also adopt the Einstein summation convention throughout. Applying Eq. (5) to a GP yields

$$f_i | x_i, \{f_{j^*}, x_{j^*}\} \sim \mathcal{N}\left(\langle f_i \rangle + K_{ik^*} (K^{-1})_{k^*j^*} (f_{j^*} - \langle f_{j^*} \rangle), K_{ij} - K_{im^*} (K^{-1})_{m^*n^*} K_{n^*j}\right). \quad (6)$$

As the set $\{f_{j^*}\}$ of known values is expanded, the conditioned processes from which the ordinates are drawn become tighter, reducing the variance between different realizations f of the GP. That is, the uncertainty in the representation decreases as knowledge of the function accrues.

Furthermore, conditioned processes can be inferred simultaneously for f_i and its first derivative $\partial_i f := df/dx|_{x=x_i}$ via a self-consistent joint distribution (see Secs. 4.1.1 and 9.4 of Ref. [61] for a derivation)

$$\begin{bmatrix} f_i \\ \partial_i f \end{bmatrix} \Big| x_i \sim \mathcal{N}\left(\begin{bmatrix} \langle f_i \rangle \\ \langle \partial_i f \rangle \end{bmatrix}, \begin{bmatrix} K_{ij} & \partial_j K_{i\bullet} \\ \partial_i K_{\bullet j} & \partial_i \partial_j K_{\bullet\bullet} \end{bmatrix}\right). \quad (7)$$

Here we introduce the notation $\partial_i K_{\bullet j} := \partial_x K(x, x_j)|_{x=x_i}$, $\partial_j K_{i\bullet} := \partial_{x'} K(x_i, x')|_{x'=x_j}$, and $\partial_i \partial_j K_{\bullet\bullet} := \partial_x \partial_{x'} K(x, x')|_{x=x_i, x'=x_j}$. Applying Eq. (5), the conditioned distribution becomes

$$\begin{bmatrix} f_i \\ \partial_i f \end{bmatrix} \Big| x_i, \{f_{j^*}, x_{j^*}\} \sim \mathcal{N}\left(\mathbb{E}\left(\begin{bmatrix} f_i \\ \partial_i f \end{bmatrix}\right), \text{Cov}\left(\begin{bmatrix} f_i \\ \partial_i f \end{bmatrix}, \begin{bmatrix} f_j \\ \partial_j f \end{bmatrix}\right)\right), \quad (8)$$

where

$$\mathbb{E}\left(\begin{bmatrix} f_i \\ \partial_i f \end{bmatrix}\right) = \begin{bmatrix} \langle f_i \rangle \\ \langle \partial_i f \rangle \end{bmatrix} + \begin{bmatrix} K_{ij^*} \\ \partial_i K_{\bullet j^*} \end{bmatrix} (K^{-1})_{j^*k^*} (f_{k^*} - \langle f_{k^*} \rangle), \quad (9)$$

$$\text{Cov}\left(\begin{bmatrix} f_i \\ \partial_i f \end{bmatrix}, \begin{bmatrix} f_j \\ \partial_j f \end{bmatrix}\right) = \begin{bmatrix} K_{ij} & \partial_j K_{i\bullet} \\ \partial_i K_{\bullet j} & \partial_i \partial_j K_{\bullet\bullet} \end{bmatrix} - \begin{bmatrix} K_{ij^*} \\ \partial_i K_{\bullet j^*} \end{bmatrix} (K^{-1})_{j^*k^*} \begin{bmatrix} K_{k^*j} & \partial_j K_{k^*\bullet} \end{bmatrix}, \quad (10)$$

given known values $\{f_{j^*}\}$ and corresponding abscissae $\{x_{j^*}\}$. Contingent on our choice of K and hyperparameters, this conditioned process allows us to make probabilistic statements about f_i and $\partial_i f$ at arbitrary abscissae based on observations of ordinates in other regions of the function's domain. In other words, Eq. (8) generates a function over its whole domain from a finite set of known values. GPR is more sophisticated than simple interpolation, as it automatically comes equipped with estimates of the uncertainty in the interpolated values.

In our application, we use GPR to model uncertainty in the NS EOS. Viewing the GP as a generator of random functions with (specifiable) common overall behavior, we harness GPR to produce many random synthetic EOSs that resemble candidate EOSs from nuclear theory. We find that the squared-exponential covariance suitably reproduces the variability observed in the nuclear-theoretic candidate EOSs when the hyperparameters are appropriately selected. In principle, one could choose the hyperparameters by optimizing the likelihood $P(\{f_{j^*}\}|\{x_{j^*}\}, \sigma, l)$ of observing the known data f_{j^*} at x_{j^*} with respect to (σ, l) ; or, better still, marginalize over the hyperparameters according to this likelihood and a prior $P(\sigma, l)$ so that

$$P(f_i | x_i, \{f_{j^*}, x_{j^*}\}) = \int d\sigma dl P(f_i | x_i, \{f_{j^*}, x_{j^*}\}; \sigma, l) P(\{f_{j^*}\} | \{x_{j^*}\}; \sigma, l) P(\sigma, l), \quad (11)$$

in effect defining a mixture model of many GPs. However, in our implementation we simply set the hyperparameters by inspection, such that the GP can reproduce the features of the candidate EOSs. This manual selection of hyperparameters does not influence our results, and exploring refinements of these choices is left to future work.

B. Gaussian process model for the equation of state

The NS EOS is a function $\varepsilon(p)$ that relates the star's pressure p to its total energy density ε , the sum of rest-mass energy density ρ and internal energy density ϵ . The rest-mass and total energy densities are related by the first law of thermodynamics,

$$d\varepsilon = \frac{(\varepsilon + p)}{\rho} d\rho. \quad (12)$$

Since the precise functional relationship between ε and p is uncertain, but partial knowledge of the nuclear microphysics places some restrictions on its behavior, a GP conditioned on candidate models from nuclear theory is especially well-suited to modeling EOS uncertainty. In particular, individual realizations of the GP serve as samples from the set of physically viable EOSs.

The simplest way to construct a non-parametric representation of the EOS would be to apply GPR directly in the p - ε plane. However, the resulting GP would formally support EOSs with $\varepsilon < 0$, since its range is necessarily the whole real line. Moreover, by default, GPR on $\varepsilon(p)$ would include thermodynamically unstable or acausal EOSs—i.e. functions with $dp/d\varepsilon < 0$ or superluminal sound speeds $c_s = \sqrt{dp/d\varepsilon} > c$. To incorporate causality, thermodynamic stability, and positivity of the total energy automatically, we instead build a GP over

$$\phi = \log \left(c^2 \frac{d\varepsilon}{dp} - 1 \right). \quad (13)$$

The auxiliary variable ϕ , first introduced in Ref. [35], can take any value along the real line, and it naturally incorporates the desired physical constraints because $\phi \in \mathbb{R}$ corresponds to $0 \leq dp/d\varepsilon \leq c^2$. Positivity of p then ensures $c_s \leq c$ and $\varepsilon \geq 0$. The spectral approach to EOS inference [36, 37] employed in Ref. [15] also starts with a similar transformation, but goes on to decompose the EOS onto a small set of basis functions. In contrast, we assign a prior process to $\phi = \phi(\log p)$ via GPR conditioned on tabulated EOSs. The process over ϕ can easily be translated to a process over ε . Since the map from ϕ to ε is nonlinear, Gaussian uncertainty in ϕ will, however, generally not correspond to Gaussian uncertainty in ε .

We condition the GP for $\phi(\log p)$ on a training set of n candidate EOSs $\varepsilon^{(\alpha)}(p)$ from the literature. The data for the α^{th} candidate EOS constitute a function $\log \varepsilon^{(\alpha)}(\log p_{j*}^{(\alpha)})$ with ordinates $\log \varepsilon_{j*}^{(\alpha)}$ and abscissae $\log p_{j*}^{(\alpha)}$.³ As the amount of $\{\log \varepsilon_{j*}^{(\alpha)}, \log p_{j*}^{(\alpha)}\}$ data available may vary from candidate EOS to candidate EOS, and the relative weight assigned to each model in the training set is proportional to the number of data points included, we resample each $\varepsilon^{(\alpha)}(p)$ to s points so that the GP for $\phi(\log p)$ is conditioned equally on every input EOS. This need not be the case—one could formulate a mixture model of GPs (see, e.g., Refs. [62, 63]) to establish a weighted training set of candidate EOSs—but for simplicity we assume equal weights here. This resampling could be performed with, e.g., linear interpolation, but we instead use GPR to obtain an estimate of the uncertainty associated with the interpolation. Thus, we construct a GP representation of $\phi^{(\alpha)}(\log p)$ for each candidate EOS in the training set, such that one realization of the α^{th} GP is an s -fold list of ordinates $\{\phi_i^{(\alpha)}\}$ at evenly spaced points in $\log p$. In this way, the maps from ε to ϕ for the tabulated EOSs—which come without uncertainties—are effectively equipped with error bars. The GPs for $\phi^{(\alpha)}(\log p)$ are subordinate to, and used as input for, the overarching GP for $\phi(\log p)$. We next describe the construction of the GPs for $\phi^{(\alpha)}(\log p)$ in some detail, before addressing the GP for $\phi(\log p)$ itself.

For every candidate EOS $\varepsilon^{(\alpha)}(p)$, we first fit $\log \varepsilon^{(\alpha)}$ with a low-order polynomial in $\log p$ and construct a GP for the residuals. The resulting joint distribution on $\log \varepsilon^{(\alpha)}$, its first derivative $\partial_i \log \varepsilon^{(\alpha)} := d \log \varepsilon^{(\alpha)} / d \log p|_{p=p_i}$ and the tabulated data $\{\log \varepsilon_{j*}^{(\alpha)}\}$ is

$$\left[\begin{array}{c} \log \varepsilon_i^{(\alpha)} \\ \partial_i \log \varepsilon^{(\alpha)} \\ \log \varepsilon_{j*}^{(\alpha)} \end{array} \right] \left| \log p_i, \{\log p_{j*}^{(\alpha)}\} \right. \sim \mathcal{N} \left(\left[\begin{array}{c} \log \hat{\varepsilon}_i^{(\alpha)} \\ \partial_i \log \hat{\varepsilon}^{(\alpha)} \\ \log \hat{\varepsilon}_{j*}^{(\alpha)} \end{array} \right], \left[\begin{array}{ccc} K_{ij}^{(\alpha)} & \partial_j K_{i*}^{(\alpha)} & K_{ij*}^{(\alpha)} \\ \partial_i K_{*j}^{(\alpha)} & \partial_i \partial_j K_{**}^{(\alpha)} & \partial_i K_{*j*}^{(\alpha)} \\ K_{i*j}^{(\alpha)} & \partial_j K_{i*}^{(\alpha)} & K_{i*j*}^{(\alpha)} \end{array} \right] \right). \quad (14)$$

Here we take the values $\{\log \hat{\varepsilon}_i^{(\alpha)}, \partial_i \log \hat{\varepsilon}^{(\alpha)}, \log \hat{\varepsilon}_{j*}^{(\alpha)}\}$ from the low-order polynomial fit as the mean in the process and model correlations among the residuals around this fit. While we assume a squared-exponential covariance kernel for all tabulated EOSs, we choose hyperparameters $(\sigma^{(\alpha)}, l^{(\alpha)})$ by hand separately for each $K_{ij}^{(\alpha)}$ and inspect the resulting GPs to ensure that they adequately reproduce the features of the input EOSs.

Conditioning the joint distribution on the tabulated ordinates $\{\log \varepsilon_{j*}^{(\alpha)}\}$ via Eq. (5), we obtain

$$\left[\begin{array}{c} \log \varepsilon_i^{(\alpha)} \\ \partial_i \log \varepsilon^{(\alpha)} \end{array} \right] \left| \log p_i, \{\log \varepsilon_{j*}^{(\alpha)}, \log p_{j*}^{(\alpha)}\} \right. \sim \mathcal{N} \left(\mathbf{E}^{(\alpha)} \left(\left[\begin{array}{c} \log \varepsilon_i \\ \partial_i \log \varepsilon \end{array} \right] \right), \text{Cov}^{(\alpha)} \left(\left[\begin{array}{c} \log \varepsilon_i \\ \partial_i \log \varepsilon \end{array} \right], \left[\begin{array}{c} \log \varepsilon_j \\ \partial_j \log \varepsilon \end{array} \right] \right) \right), \quad (15)$$

³ Candidate EOS data tabulated as p vs. ρ can also be accommodated by transforming the rest-mass energy density to total energy density via Eq. (12).

where the explicit expressions for the expectation value and conditioned covariance follow from Eqs. (9) and (10). The corresponding process for $\phi^{(\alpha)}$ is calculated from Eq. (13) as

$$\phi_i^{(\alpha)} \mid \log p_i, \left\{ \log \varepsilon_{j*}^{(\alpha)}, \log p_{j*}^{(\alpha)} \right\} \sim \mathcal{N} \left(\mathbf{E}^{(\alpha)}(\phi_i), \text{Cov}^{(\alpha)}(\phi_i, \phi_j) \right), \quad (16)$$

with

$$\mathbf{E}^{(\alpha)}(\phi_i) = \log \left(\mathbf{E}^{(\alpha)}(\partial_i \log \varepsilon) \left(\frac{e^{\mathbf{E}^{(\alpha)}(\log \varepsilon_i)}}{p_i} c^2 - 1 \right) \right). \quad (17)$$

We approximate the covariance matrix $\text{Cov}^{(\alpha)}(\phi_i, \phi_j)$ through a first order Taylor expansion for ϕ_i in terms of $\log \varepsilon_i$ and $\partial_i \log \varepsilon$, i.e.

$$\phi_i(\log \varepsilon_i, \partial_i \log \varepsilon) \approx \phi_i + \left. \frac{\delta \phi}{\delta \log \varepsilon} \right|_i \delta \log \varepsilon_i + \left. \frac{\delta \phi}{\delta(\partial \log \varepsilon)} \right|_i \delta(\partial_i \log \varepsilon), \quad (18)$$

so that

$$\begin{aligned} \text{Cov}(\phi_i, \phi_j) &= \left(\left. \frac{\delta \phi}{\delta \log \varepsilon} \right|_i \right) \left(\left. \frac{\delta \phi}{\delta \log \varepsilon} \right|_j \right) \text{Cov}(\log \varepsilon_i, \log \varepsilon_j) \\ &+ \left(\left. \frac{\delta \phi}{\delta(\partial \log \varepsilon)} \right|_i \right) \left(\left. \frac{\delta \phi}{\delta \log \varepsilon} \right|_j \right) \text{Cov}(\partial_i \log \varepsilon, \log \varepsilon_j) \\ &+ \left(\left. \frac{\delta \phi}{\delta \log \varepsilon} \right|_i \right) \left(\left. \frac{\delta \phi}{\delta(\partial \log \varepsilon)} \right|_j \right) \text{Cov}(\log \varepsilon_i, \partial_j \log \varepsilon) \\ &+ \left(\left. \frac{\delta \phi}{\delta(\partial \log \varepsilon)} \right|_i \right) \left(\left. \frac{\delta \phi}{\delta(\partial \log \varepsilon)} \right|_j \right) \text{Cov}(\partial_i \log \varepsilon, \partial_j \log \varepsilon). \end{aligned} \quad (19)$$

This allows us to translate $\log \varepsilon^{(\alpha)}(\log p_{j*}^{(\alpha)})$ to the auxiliary variable $\phi^{(\alpha)}(\log p_i)$, interpolated at an arbitrary set of abscissae, while keeping track of the uncertainty associated with the interpolation. We refer to this as *resampling* the tabulated EOSs, because we take irregularly spaced tabulated data for $\varepsilon^{(\alpha)}(p)$ and transform it into a process for $\phi^{(\alpha)}(\log p)$ from which we can extract regularly sampled data. Fig. 1 shows the $\phi^{(\alpha)}(\log p)$ process we construct for one of the candidate EOSs described in Sec. IV.

The overarching GP for $\phi(\log p)$ is subsequently conditioned on the collection of processes $\phi^{(\alpha)}(\log p)$ for the individual candidate EOSs. We again use a squared-exponential kernel, but supplement it with a white-noise kernel (cf. Eq. (4))

$$K_{ij}^{(\text{EOS})} = \delta_{ij} \frac{\sigma_n^2}{n} \sum_{\alpha} \left(\mathbf{E}^{(\alpha)}(\phi_i) - \mathbf{E}^{(\text{EOS})}(\phi_i) \right)^2 \quad (20)$$

scaled by the observed variance between the n tabulated EOSs at the pressures p_i , where

$$\mathbf{E}^{(\text{EOS})}(\phi_i) = \frac{1}{n} \sum_{\beta} \mathbf{E}^{(\beta)}(\phi_i). \quad (21)$$

In this way, we not only model the covariance between ϕ at different pressures, but also the spread in candidate EOSs at each abscissa. The scaling hyperparameter σ_n in Eq. (20) allows us to specify how closely we wish the synthetic EOSs to follow the tabulated ones; small values enforce the conditioned process to closely follow the average of the tabulated EOSs. For simplicity, we condition based on the interpolated ordinates $\phi_{i*}^{(\alpha)}$ at the same regularly sampled set of pressures $\log p_{i*}$ regardless of the EOS, though this is not strictly necessary.

Equipped with processes $\phi^{(\alpha)}(\log p)$ for each tabulated EOS, and the white-noise kernel approximating modeling uncertainty at each pressure, we construct the overarching GP for $\phi(\log p)$ conditioned on all the tabulated EOSs. The joint distribution on $\phi(\log p)$ and the subordinate $\phi^{(\alpha)}(\log p)$ is thus

$$\left[\begin{array}{c} \phi_k \\ \phi_{i*}^{(\alpha=1)} \\ \phi_{i*}^{(\alpha=2)} \\ \vdots \\ \phi_{i*}^{(\alpha=n)} \end{array} \right] \left| \log p_k, \log p_{i*} \sim \mathcal{N}(\mathbf{E}, \text{Cov}) \quad (22)$$




FIG. 1. A demonstration of our GP mapping from $\varepsilon^{(h4)}$ to $\phi^{(h4)}$ using a 3rd-order polynomial fit and a squared-exponential kernel with $\sigma^{(h4)} = 0.1$ and $l^{(h4)} = 0.6$ (see Eq. (14)). Irregularly sampled tabulated values and estimates obtained from numeric differentiation (*red*) are resampled to a regular grid in $\log p$ via GPR, along with residuals and associated error estimates (*gray*). We note that the residuals between tabulated values and our interpolation of $\varepsilon^{(\alpha)}$ are $\lesssim 1\%$ at high pressures, resulting in $\lesssim 10\%$ relative uncertainty in ϕ .

with

$$\mathbf{E} = \begin{bmatrix} \hat{\mathbf{E}}^{(\text{EOS})}(\phi_k) \\ \hat{\mathbf{E}}^{(\text{EOS})}(\phi_{i^*}) \\ \hat{\mathbf{E}}^{(\text{EOS})}(\phi_{i^*}) \\ \vdots \\ \hat{\mathbf{E}}^{(\text{EOS})}(\phi_{i^*}) \end{bmatrix}, \quad (23)$$

where $\hat{\mathbf{E}}^{(\text{EOS})}(\phi_{i^*})$ is a low-order polynomial fit to $\mathbf{E}^{(\text{EOS})}(\phi_{i^*})$ in $\log p$, and

$$\text{Cov} = \begin{bmatrix} K_{ij} & K_{ij^*} & K_{ij^*} & \cdots & K_{ij^*} \\ K_{i^*j} & K_{i^*j^*} + \text{Cov}_{i^*j^*}^{(\alpha=1)} + K_{i^*j^*}^{(\text{EOS})} & K_{i^*j^*} & \cdots & K_{i^*j^*} \\ K_{i^*j} & K_{i^*j^*} & K_{i^*j^*} + \text{Cov}_{i^*j^*}^{(\alpha=2)} + K_{i^*j^*}^{(\text{EOS})} & \cdots & K_{i^*j^*} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ K_{i^*j} & K_{i^*j^*} & K_{i^*j^*} & \cdots & K_{i^*j^*} + \text{Cov}_{i^*j^*}^{(\alpha=n)} + K_{i^*j^*}^{(\text{EOS})} \end{bmatrix}. \quad (24)$$

The covariance matrix for this GP representation of the NS EOS is composed of several terms:

- block diagonal contributions $\text{Cov}^{(\alpha)}$ from the conditioned uncertainty in $\phi_i^{(\alpha)}$ for each candidate EOS, as in Eq. (16);
- diagonal contributions $K^{(\text{EOS})}$ from the white-noise kernel (20), encoding modeling uncertainty in the ordinates $\{\phi_{j^*}^{(\alpha)}\}$ from the candidate EOSs at each pressure p_{j^*} ; and
- a squared-exponential covariance matrix K , as in Eq. (3), describing the correlations between all ordinates ϕ_i .

Specific choices for the squared-exponential kernel hyperparameters (σ , l) and the white-noise kernel scaling factor σ_n complete the model. Conditioning this process on the means $\mathbf{E}^{(\alpha)}$ for the tabulated EOSs yields

$$\phi_i | \log p_i, \{\mathbf{E}^{(\alpha=1)}(\phi_{j^*}), \mathbf{E}^{(\alpha=2)}(\phi_{j^*}), \dots, \mathbf{E}^{(\alpha=n)}(\phi_{j^*}), \log p_{j^*}\} \sim \mathcal{N}(\tilde{\mathbf{E}}, \tilde{\text{Cov}}), \quad (25)$$




FIG. 2. Inferred processes for several $\phi^{(\alpha)}(\log p)$ (red) and an example overarching GP $\phi(\log p)$ (black) conditioned on all of them (left) along with the same reference EOSs and the model-agnostic (cyan, loosely informed by nuclear theory) and model-informed (blue, more tightly constrained by nuclear theory) priors used in this analysis (right, see Sec. IV). Shaded regions correspond to 1- σ confidence regions. We see that the model-agnostic prior comfortably contains all the reference EOSs, while the model-informed prior more closely follows a group of tabulated EOS.

which approximates the desired process $\phi_i | \log p_i, \{\log \varepsilon_{j_*}^{(\alpha=1)}, \log p_{j_*}^{(\alpha=1)}\}, \{\log \varepsilon_{k_*}^{(\alpha=2)}, \log p_{k_*}^{(\alpha=2)}\}, \dots, \{\log \varepsilon_{l_*}^{(\alpha=n)}, \log p_{l_*}^{(\alpha=n)}\}$ by conditioning on interpolated values $E^{(\alpha)}$ with associated uncertainties instead of conditioning directly on the tabulated EOS data. This is illustrated in Fig. 2. Concrete expressions for the final conditioned expectation value and covariance follow from Eq. (6).

The GP defined by Eq. (25) models $\phi(\log p)$, rather than $\varepsilon(p)$. Nonetheless, for each realization $\phi^{(a)}(\log p)$ of the process, we invert Eq. (13) and numerically integrate $\partial \varepsilon^{(a)} / \partial p = (1 + e^{\phi^{(a)}}) / c^2$ to obtain $\varepsilon^{(a)}(p)$, which is suitable for calculating macroscopic NS properties. The rest-mass density can be obtained from $\varepsilon^{(a)}$ via Eq. (12), allowing us to plot the EOS in conventional $p^{(a)}(\rho)$ form. The individual realizations of the EOS are stitched onto a model for the low-density crust EOS, which is well known below nuclear saturation density; we use a piecewise polytrope implementation of `sly` [38]. In this way, we obtain a *unified EOS* valid at all densities up to a specified maximum value. The distribution of functions $\varepsilon^{(a)}(p)$ calculated in this manner defines the *EOS prior process* $P(\varepsilon(p) | \{\varepsilon_{j_*}^{(\alpha)}\}, \vec{\sigma})$. Each *synthetic EOS* drawn from this process is a random function that resembles the input tabulated EOSs both qualitatively and quantitatively. The choice of hyperparameters $\vec{\sigma} := [\sigma, l, \sigma_n, \sigma^{(\alpha=1)}, l^{(\alpha=1)}, \sigma^{(\alpha=2)}, l^{(\alpha=2)}, \dots, \sigma^{(\alpha=n)}, l^{(\alpha=n)}]$ for the GP provides the freedom to self-consistently explore a wide swath of the space of possible EOSs, or to mimic to the behavior observed in the candidate EOSs more closely.

The process for $\varepsilon(p)$, conditioned on the full training set of candidate EOSs, generates a large set of new proposals for the pressure-density relation without resorting to a parameterization. Astrophysical constraints on the EOS, such as the requirement that it support observed NS masses, are applied by directly checking, e.g., the maximum mass for each synthetic EOS. Any synthetic EOS that violates the constraint is discarded from the prior. Similar astrophysical cutoffs could incorporate future information about NS radii and moments of inertia from electromagnetic astronomy.

III. BAYESIAN INFERENCE OF THE EQUATION OF STATE VIA MONTE CARLO INTEGRATION

GP representations of the uncertainty in NS internal structure play a fundamental role in our non-parametric inference of the EOS. The candidate EOSs upon which the GP is conditioned, as well as the choice of the hyperparameters $\vec{\sigma}$, determine the characteristics of the EOS prior process $P(\varepsilon(p) | \{\varepsilon_{j_*}^{(\alpha)}\}, \vec{\sigma})$. In particular, the degree to which elements $\varepsilon^{(a)}(p) \sim P(\varepsilon(p) | \{\varepsilon_{j_*}^{(\alpha)}\}, \vec{\sigma})$ of the prior process resemble the input EOSs is controlled by the hyperparameters σ, l

and σ_n . Here we show how the GW likelihood from parameter estimation is used in conjunction with the prior to produce an *EOS posterior process*.

The posterior process for the EOS is ultimately derived from GW detector data (d), which is sensitive to macroscopic NS observables like masses and tidal deformabilities. Source properties are recovered from d via Bayesian inference for the parameters of a prescribed waveform model. The waveform parameters include the masses $M_{1,2}$, the tidal deformabilities $\Lambda_{1,2}$, and other source properties and phenomenological parameters denoted as ϑ . The priors chosen for $M_{1,2}$ and ϑ , along with the underlying description of the waveform, constitute the Bayesian model \mathcal{H} for the data.

Strictly speaking, the waveform parameters $M_{1,2}$ and $\Lambda_{1,2}$ are not completely independent, as a specification of M and the EOS is sufficient to determine Λ . Indeed, an EOS $\varepsilon(p)$ predicts a unique M - Λ relation, which we denote as $\Lambda(M; \varepsilon)$. This relation is calculated by first integrating the Tolman-Oppenheimer-Volkoff (TOV) equations [64, 65] and the field equation governing the quadrupolar tidal perturbation [66] simultaneously from the center of the NS to its surface for a choice of central density ρ_c . These equations of stellar structure determine the mass, radius and tidal deformability of the NS. A sequence of stable neutron stars is then constructed by repeating the integration for a sequence of central densities up to ρ_c^{\max} , the central density for which the mass reaches a maximum M_{\max} and beyond which the star becomes unstable [67]. The resulting tidal deformability and mass sequence constitutes $\Lambda(M; \varepsilon)$. Integration of the equations of stellar structure also gives us access to the M - ρ_c relation $M(\rho_c; \varepsilon)$, which includes M_{\max} information.

The EOS posterior process $P(\varepsilon|d; \{\varepsilon_{j*}^{(\alpha)}\}, \vec{\sigma}, \mathcal{H})$ is obtained from the GW data through an application of Bayes' theorem, which states

$$P(\varepsilon|d; \{\varepsilon_{j*}^{(\alpha)}\}, \vec{\sigma}, \mathcal{H}) \propto P(d|\varepsilon; \mathcal{H}) P(\varepsilon|\{\varepsilon_{j*}^{(\alpha)}\}, \vec{\sigma}). \quad (26)$$

$P(d|\varepsilon; \mathcal{H})$ is the likelihood of the data given an EOS $\varepsilon(p)$ and the model \mathcal{H} . Marginalizing separately over the waveform parameters of interest ($M_{1,2}, \Lambda_{1,2}$) and the nuisance parameters ϑ , it can be written as

$$P(d|\varepsilon; \mathcal{H}) = \int dM_1 dM_2 P(M_1, M_2|\mathcal{H}) \left[\int d\vartheta P(\vartheta|\mathcal{H}) P(d|\vartheta, M_1, M_2, \Lambda_1(M_1; \varepsilon), \Lambda_2(M_2; \varepsilon)) \right]. \quad (27)$$

The integral in square brackets is the marginal conditional likelihood $\mathcal{L}(d|M_1, M_2, \Lambda_1(M_1; \varepsilon), \Lambda_2(M_2; \varepsilon); \mathcal{H})$ for the observed GW data given a pair of masses and tidal deformabilities. In terms of this likelihood, Eq. (26) reads

$$\begin{aligned} P(\varepsilon_i|d; \{\varepsilon_{j*}^{(\alpha)}\}, \vec{\sigma}, \mathcal{H}) &\propto P(\varepsilon_i|\{\varepsilon_{j*}^{(\alpha)}\}, \vec{\sigma}) \\ &\times \int d\Lambda_1 d\Lambda_2 dM_1 dM_2 P(M_1, M_2|\mathcal{H}) \delta(\Lambda_1 - \Lambda(M_1; \varepsilon_i)) \delta(\Lambda_2 - \Lambda(M_2; \varepsilon_i)) [\mathcal{L}(d|M_1, M_2, \Lambda_1, \Lambda_2; \mathcal{H})]. \end{aligned} \quad (28)$$

Here, we marginalize over $\Lambda_{1,2}$ in addition to $M_{1,2}$ by introducing Dirac delta functions over $\Lambda(M; \varepsilon)$. This accounts for the fact that the likelihood available from standard parameter estimation does not include intrinsic correlations among $\Lambda_{1,2}$ (cf. Ref. [28]); most stochastic samplers employed in GW parameter estimation sample from $\mathcal{L}(d|M_1, M_2, \Lambda_1, \Lambda_2; \mathcal{H})$, treating $\Lambda_{1,2}$ as completely independent parameters.⁴

Thus, provided that $\mathcal{L}(d|M_1, M_2, \Lambda_1, \Lambda_2; \mathcal{H})$ is known, we need only compute the integral (28) to obtain a posterior process for the EOS. In practice, however, standard parameter estimation gives us access to a finite number of samples from the likelihood instead of the distribution itself. Assuming that the exact marginal conditional likelihood can be accurately modeled from the available samples [68], we simply approximate $\mathcal{L}(d|M_1, M_2, \Lambda_1, \Lambda_2; \mathcal{H})$ with a Gaussian KDE (see, e.g., Ref. [28]). The integral is then evaluated with Monte Carlo techniques. Although Monte Carlo integration is relatively inefficient, it provides the most straightforward way to compute the posterior process. More advanced methods, such as Markov-Chain Monte Carlo, may be amenable to the inclusion of a GP prior, but the associated jump proposals are non-trivial (see, e.g., Refs. [69, 70]). We therefore rely exclusively on direct Monte Carlo integration in this paper, implementing the following algorithm to calculate the posterior process:

- a) draw a synthetic EOS $\varepsilon_i^{(a)} \sim P(\varepsilon_i|\{\varepsilon_{j*}^{(\alpha)}\}, \vec{\sigma})$ as a realization of the GP prior;
- b) draw a pair of component masses $M_1^{(a)}, M_2^{(a)} \sim P(M_1, M_2|\mathcal{H})$ from a specified prior;

⁴ Some parameter estimation analyses may include non-trivial priors for $M_{1,2}$ and $\Lambda_{1,2}$, but we can account for this by appropriately re-weighting their samples.

- c) compute tidal deformabilities $\Lambda_1^{(a)}$ and $\Lambda_2^{(a)}$ associated with $M_1^{(a)}$, $M_2^{(a)}$, and $\varepsilon_i^{(a)}$ via the equations of stellar structure;
- d) evaluate the corresponding GW likelihood $\mathcal{L}^{(a)} = \mathcal{L}(d|M_1^{(a)}, M_2^{(a)}, \Lambda_1^{(a)}, \Lambda_2^{(a)}; \mathcal{H})$ via the KDE; and
- e) repeat until a sufficiently large number of samples have been collected.

Because our samples have unequal weights, we gauge the precision of the Monte Carlo integral through the *effective* number of samples

$$N_{\text{eff}} = \exp \left(- \sum_a w^{(a)} \log w^{(a)} \right) \quad (29)$$

with

$$w^{(a)} = \left(\mathcal{L}^{(a)} / \sum_b \mathcal{L}^{(b)} \right) \quad (30)$$

based on the entropy of the weights observed *a posteriori*. The larger N_{eff} , the better the precision. If only a few samples have significant weights, then both the entropy and N_{eff} will be small. However, if most samples have similar weights, the entropy and N_{eff} will be large, possibly approaching the upper bound $N_{\text{eff}} = N$ achieved for equal weights.

Monte Carlo integration also provides a means of directly estimating the evidence $P(d|\{\varepsilon_{j*}^{(\alpha)}\}, \vec{\sigma}, \mathcal{H})$ for a given prior process as a functional integral over $\varepsilon(p)$:

$$\begin{aligned} P(d|\{\varepsilon_{j*}^{(\alpha)}\}, \vec{\sigma}, \mathcal{H}) &= \int \mathcal{D}\varepsilon dM_1 dM_2 P(\varepsilon|\{\varepsilon_{j*}^{(\alpha)}\}, \vec{\sigma}) P(M_1, M_2|\mathcal{H}) \mathcal{L}(d|M_1, M_2, \Lambda(M_1; \varepsilon), \Lambda(M_2; \varepsilon); \mathcal{H}) \\ &\approx \frac{1}{N} \sum_a \mathcal{L}^{(a)} \quad \left| \quad \begin{array}{l} \varepsilon^{(a)} \sim P(\varepsilon|\{\varepsilon_{j*}^{(\alpha)}\}, \vec{\sigma}) \\ M_1^{(a)}, M_2^{(a)} \sim P(M_1, M_2|\mathcal{H}) \end{array} \right. \end{aligned} \quad (31)$$

Unlike the posterior process, the sampling uncertainty in $P(d|\{\varepsilon_{j*}^{(\alpha)}\}, \vec{\sigma}, \mathcal{H})$ scales inversely with \sqrt{N} rather than $\sqrt{N_{\text{eff}}}$, although $N_{\text{eff}} \propto N$ as $N \rightarrow \infty$. Because our KDE approximation of the likelihood has an unknown normalization factor, only relative measures of evidence are meaningful, and we therefore compute the Bayes factor—the ratio of evidences—when comparing support for different priors.

IV. GAUSSIAN PROCESS EQUATION-OF-STATE PRIORS

In our specific implementation, we condition our GP EOS prior on a fiducial set of candidate EOSs selected from extant nuclear-theoretic models. Our training set of candidate EOSs is composed of seven well-established models [71]: **alf2** [72], **eng** [73], **h4** [74], **mpa1** [75], **ms1**, **ms1b** [76], and **sly** [77]. These models span a wide range of stiffnesses and support $1.93 M_{\odot}$ stars, a conservative observational bound on the maximum NS mass [78]. Our candidate EOSs are purely hadronic, except for **alf2** and **H4**, which contain color-flavor-locked quarks and hyperons, respectively.

First off, we seek to build an uninformative EOS prior process that uniformly covers the entire range of physically viable EOSs, from the stiffest possible (saturating the causality constraint) to the softest possible (marginally failing to support a $1.93 M_{\odot}$ NS). Within this range, synthetic EOSs should vary as freely as possible subject to the physical constraints discussed in Sec. II B. We therefore use the tabulated EOSs to provide only weak, general guidance to the form of the synthetic EOSs. This is achieved by setting a large white-noise variance σ_n , orders of magnitude larger than the correlation amplitude σ , while choosing σ and l of order unity. Thus, synthetic EOSs drawn from the prior can depart significantly from the mean of the tabulated EOSs upon which they are conditioned (cf. Eq. (23)) and can exhibit different types of functional behavior (cf. Eq. (24)) while exploring a large swath of the pressure-density plane. We refer to this fiducial EOS prior process as the *model-agnostic prior*.

We also explore an alternate choice of prior. Our second prior seeks to conform more closely to our tabulated EOSs, producing much tighter *a priori* bounds in the pressure-density plane. In this case, hyperparameters are chosen to yield synthetic EOS behavior that reproduces the features common to all of the input EOSs: σ_n is set to be comparable to σ , while σ and l are essentially unchanged from the model-agnostic case, and the synthetic EOSs consequently depart relatively little from the average of the candidate EOSs. We call the resulting EOS prior process the *model-informed prior*.

While the model agnostic prior approximates an uninformative process over the EOS, the model-informed prior preferentially encodes features that are common to most nuclear-theoretic models, producing a non-trivial process that favors certain candidate EOSs. Accordingly, EOSs that are outliers in, e.g., stiffness are relatively unlikely given the model-informed prior, and hence particularly stiff or soft EOSs are disfavored *a posteriori* as well as *a priori*. We note that our specific hyperparameter selections for the model-informed prior are somewhat *ad hoc* and certainly not unique. Different choices would produce different non-trivial priors, and these may be worth exploring in the future.

Our GP priors are shown in Fig. 2 along with our set of tabulated EOSs. We see that the model-agnostic prior supports a much broader range of possible EOSs, while remaining centered—like the model-informed prior—on our representative set of candidate models. We also compare our prior processes to the spectral method’s [15] in Fig. 3. One observes that the spectral prior is centered on a similar region of the pressure-density plane as both our model-agnostic and model-informed priors, but that our model-agnostic prior covers a somewhat broader region. Because our synthetic EOSs are not restricted to a single parametric form, they also encapsulate a wider range of functional behaviors than the spectral EOSs, though this may not be apparent from the two-dimensional projection in the figures.

We first evaluate the non-parametric method’s performance by analyzing simulated GW signals with both the model-agnostic and model-informed GP prior processes (Sec. V). We then place constraints on the true NS EOS (Sec. VI) using publicly available GW170817 M - Λ posteriors [68].⁵ When performing the Monte Carlo integration, we draw component masses from a uniform prior consistent with the low-spin prior of Ref. [13], and we iterate until $O(10^3)$ effective samples have been returned, typically corresponding to $O(10^5)$ draws from the prior. The two parallel inferences demonstrate how the non-parametric method can be tuned to assume different levels of *a priori* confidence in the fiducial candidate EOSs. We remark, however, that neither of the priors constructed here should be construed as an optimal representation of the variability in the NS EOS. Our selection of hyperparameters and training EOSs simply provides a qualitative illustration of possible designs for the GP prior process; the optimization of the EOS prior will be pursued in future work. For example, the fiducial set of candidate EOSs could be expanded to include more recently developed models with broader phenomenology, including first-order phase transitions or strange quark matter. The hyperparameters could also be selected more judiciously to produce an EOS prior process that covers an even larger region of the pressure-density plane; this is particularly relevant for the model-informed prior, as excessively stiff or soft EOSs tend to fall toward its edge. Nonetheless, the GP priors introduced here are sufficiently broad to demonstrate the non-parametric method’s efficacy, and we expect the model-agnostic prior in particular to resemble more carefully designed EOS prior processes.

V. SIMULATED GRAVITATIONAL WAVE SIGNALS

We perform an injection campaign to demonstrate the non-parametric method’s ability to infer a known EOS from GW data. This proof of principle study is conducted with the intention of applying the same analysis to GW170817, and, accordingly, our simulated signals are designed to resemble that event. We define three different sources with parameters similar to those inferred for GW170817, and inject the corresponding GW signals into real detector noise surrounding the event [13, 79]. In particular, the injections have component masses and spins that are compatible with the inferred source properties for GW170817; the sources are placed at a comparable distance and orientation relative to the detectors; and the signals’ network signal-to-noise ratios are approximately the same ($\text{SNR} \approx 35$) [28]. The injected component spins are restricted to $\chi < 0.1$; while spins can impact the estimation of tidal effects, those inferred for GW170817 are sufficiently small that they are not thought to significantly affect our analysis. Nonetheless, while we assume the spins are small, we do not fix them to be identically zero, and we marginalize over spin when recovering the injected parameters. For each of these sources, we further select one of three candidate EOSs—**ms1**, **h4**, **sly**—ranging from relatively stiff to soft. The choice of injected masses and EOS determines the tidal deformabilities in the simulated GW signal. The $M_{1,2}$ and $\Lambda_{1,2}$ values for each mock event, as well as the associated EOS, are listed in Table I. Injections are performed with the **TaylorF2** waveform (see, e.g., Ref. [80]) and the $M_{1,2}$ - $\Lambda_{1,2}$ marginal likelihoods are sampled with **LALInference** [24, 49] using the same.

The posteriors obtained for one simulated event, injection I with **h4**, are shown by way of example in Figs. 4 and 5 along with the associated priors. Fig. 4 shows projections of the joint probability distribution over $M_{1,2}$ and $\Lambda_{1,2}$ into the Λ_1 - Λ_2 and Λ_1 - M_1 subspaces, as well as marginal distributions for M_1 , Λ_1 and Λ_2 . On the left, one can see that the model-agnostic prior supports a wide range of possible M - Λ relations, and the inferred posterior comfortably includes the injected value. This is in keeping with our expectations, since the model-agnostic prior is designed to contain all feasible EOSs. Nonetheless, the inference produced informative constraints on the tidal deformabilities, as the 90% confidence region of the posterior is much reduced relative to the prior.

⁵ Our analysis technically requires M - Λ *likelihoods*, but because Ref. [28] assumed flat priors, the likelihood is proportional to the posterior distribution for these variables.

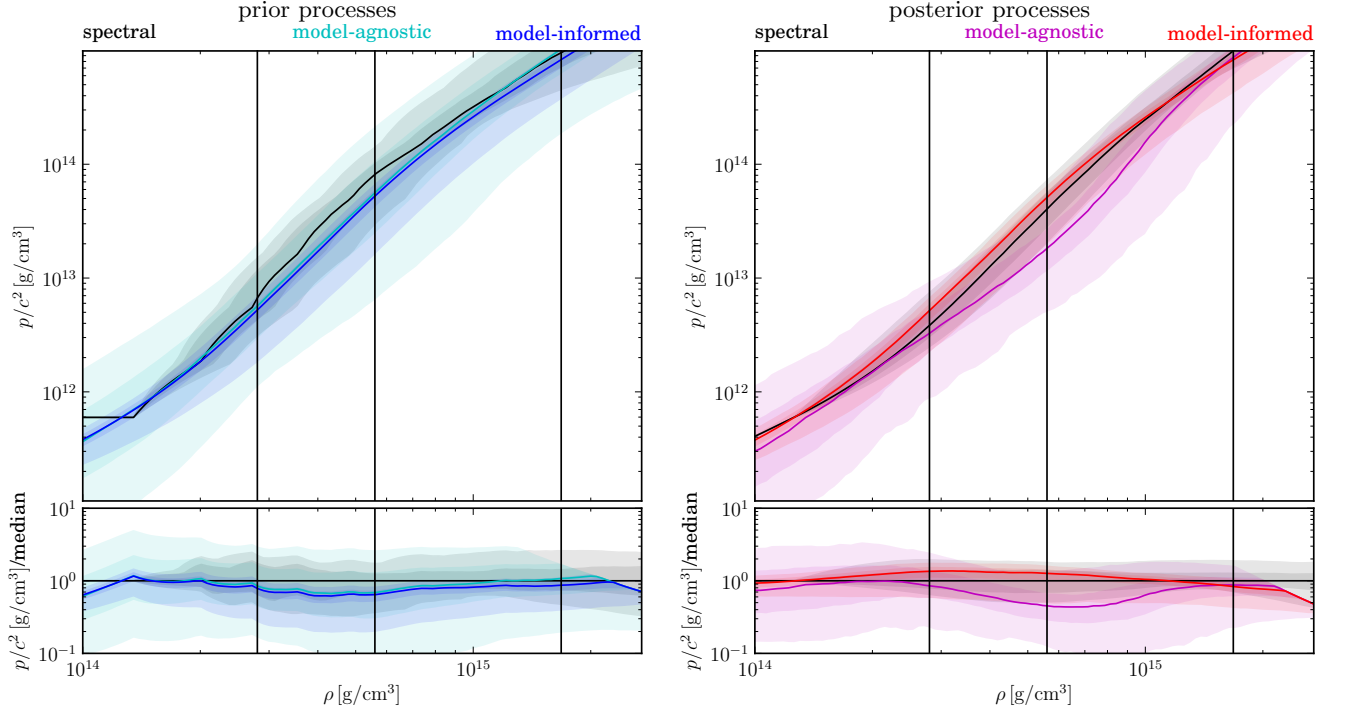


FIG. 3. A comparison for GW170817 between the spectral prior from Ref. [15] (*black*) and our model-agnostic (*cyan*) and model-informed (*blue*) priors (*left*), as well as between the spectral posteriors (*black*) and our model-agnostic (*magenta*) and model-informed (*red*) posteriors (*right*). Solid lines show the medians and shaded regions correspond to 50% and 90% credible regions. We see that our model-agnostic posterior favors even softer EOSs than the spectral approach, whereas the model-informed prior favors slightly stiffer EOSs than the spectral approach. However, the associated uncertainties suggest all results are consistent.

Fig. 5 shows the corresponding EOS prior and posterior processes in the pressure-density plane. In the lefthand panel, the broad model-agnostic prior is significantly narrowed by the inference, especially around $2\rho_{\text{nuc}}$, producing a posterior process that favors an EOS somewhat softer than the prior’s median. The inference also yields information about the central density ρ_c and pressure p_c of each NS, since the M - ρ_c relation for each synthetic EOS is known. The overplotted contours indicate the posterior for the more massive component’s central values. Constraints on the EOS above $\rho_{c,1}$, the central density associated with M_1 , arise exclusively through correlations within the synthetic EOSs; the GW likelihood directly constrains the EOS only below $\rho_{c,1}$.

The righthand panels of Figs. 4 and 5 show our inference with the model-informed prior. In terms of the joint $M_{1,2}$ - $\Lambda_{1,2}$ probability distribution, the model-informed prior is much more restrictive, leading to a tighter posterior than in the model-agnostic case. Specifically, the model-informed posterior has limited support for very soft EOSs (small tidal deformabilities). This is consistent with the fact that the model-informed prior is closely tied to the training set of candidate EOSs, which includes models no softer than **sly** ($\Lambda_{1,4} \approx 290$), whereas the less constrained synthetic EOSs from the model-agnostic prior explore softer regions of the pressure-density plane. Despite the model-informed posterior’s smaller support, the injected parameters $M_{1,2}$ and $\Lambda_{1,2}$ are still recovered within 90% confidence, although $\tilde{\Lambda}$ happens to lie marginally outside the 90% confidence region for this event.

In the model-informed case, we note that the posterior in the pressure-density plane is nearly identical to the prior. This is because the model-informed prior is already quite narrow, and a single GW event conveys only a limited amount of information. We expect tighter constraints to be achievable with combined observations from multiple events [39, 42].

The full set of injected signals is presented in Appendix A, and we summarize the results here, reporting our method’s performance for each of the injections in Table I. We quote the injected $M_{1,2}$, $\Lambda_{1,2}$ and $\tilde{\Lambda}$ values along with one-dimensional 90% credible regions around the maxima *a posteriori*. In general, we find that the component masses are always well recovered with both priors, while $\Lambda_{1,2}$ are typically well recovered with the model-agnostic prior, although the maxima *a posteriori* are not always centered on the injected values. This is due to a combination of noise fluctuations (the likelihood does not peak at the injected values) and the influence of our priors (the injected values lie near the edge of the prior). In other words, if the injected EOS is rare *a priori*, as is the case for some of the candidate EOSs in the model-informed prior, it is not surprising that they are also rare *a posteriori*. We expect

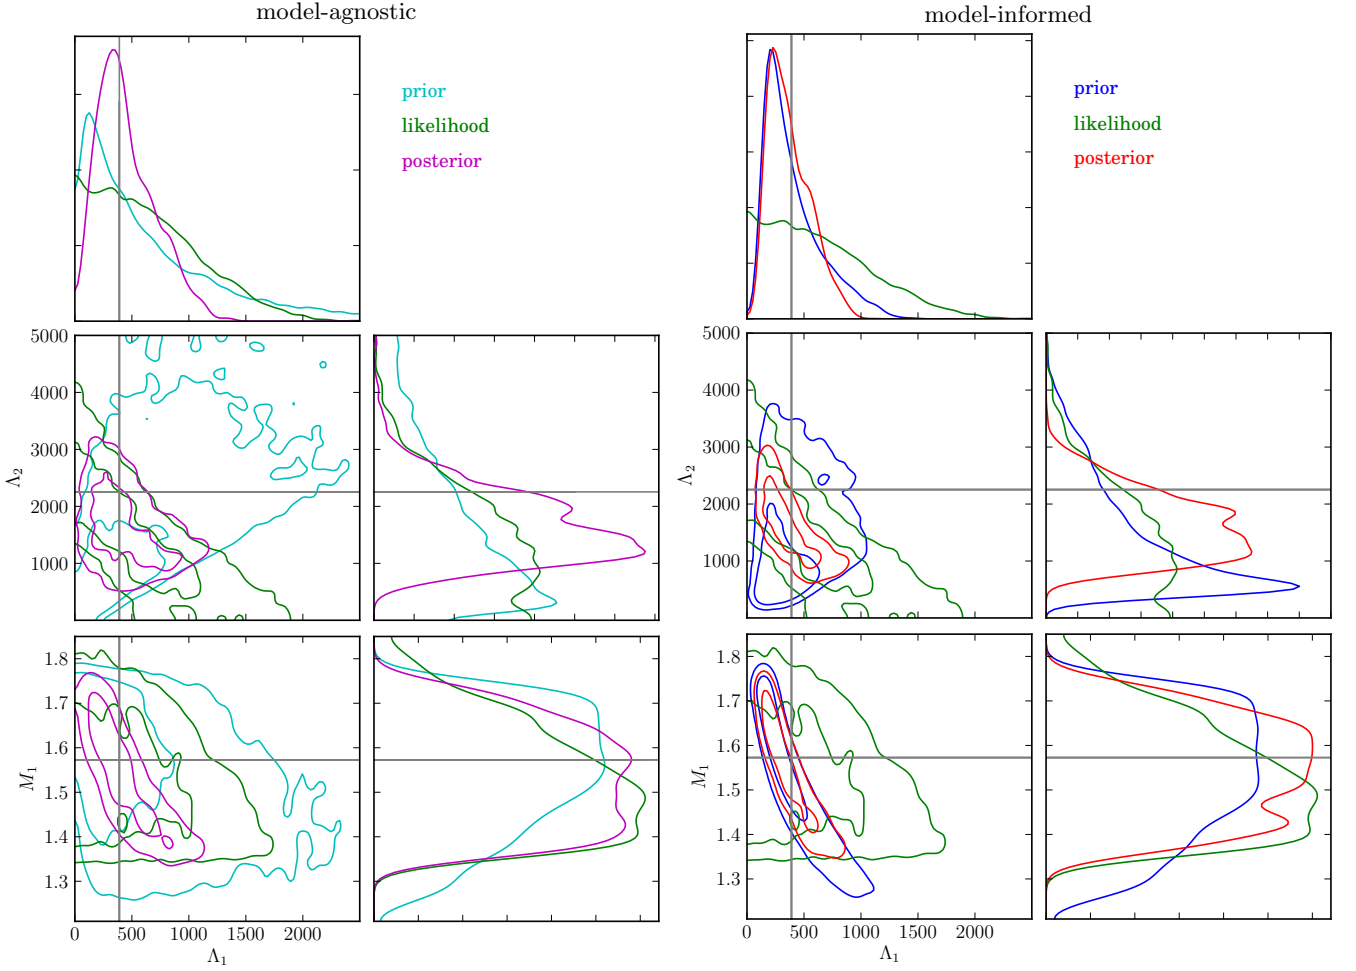


FIG. 4. Prior (cyan/blue), likelihood (green), and posterior (magenta/red) distributions for our model-agnostic (left) and model-informed (right) GP priors for a single simulated injection, event I with **h4** from Table I. Contours for the joint distributions correspond to 50% and 90% credible regions. Because of the broader support in the model-agnostic prior, the injected value (gray cross-hairs) is comfortably contained in the associated posterior’s support. We note that the model-informed prior appears to contain the injected value in some marginal distributions, but it actually lives near the edge of the allowed M - Λ region. This accounts for the fact that the $\Lambda_{1,2}$ posterior barely contains the injection even though it appears to be contained in the $\Lambda_{1,2}$ prior. Since we perform inference with a non-trivial prior in a four-dimensional space, the resulting constraints may not be obvious from two-dimensional projections.

the model-agnostic prior to recover injected EOSs more successfully because it is essentially uninformative, whereas the model-informed prior is non-trivial by design. Unsurprisingly, $\tilde{\Lambda}$ is generally less sensitive to the priors and is consequently recovered more robustly. The GW data constrain $\Lambda_{1,2}$ primarily through $\tilde{\Lambda}$, so *a posteriori* relative degrees of belief for combinations of $\Lambda_{1,2}$ that produce the same $\tilde{\Lambda}$ are mostly informed by the prior.

In contrast to the model-agnostic prior, $\Lambda_{1,2}$ are not always well recovered with the model-informed prior. While several two-dimensional projections, such as the marginal Λ_1 - Λ_2 distribution, appear to comfortably contain the injected signal, in reality it lies near the edge of the allowed M - Λ relations in some cases (see, e.g., Fig. 4). Hence, the injection is found near the edge of the posterior because of the constraints applied in the full four-dimensional space. This indicates that our model-informed prior could be too tight to reliably infer the true EOS if it is significantly softer or stiffer than the candidate EOSs included in our training set. Our model-agnostic prior is more trustworthy in this regard, though it generally produces looser constraints because of its broader support. What’s more, the model-informed prior tends to more strongly influence the inferred $\Lambda_{1,2}$, resulting at times in injected values recovered outside the 90% credible regions. However, this is to be expected for injected EOSs lying near or beyond the edge of our prior. We only expect our confidence regions to have the correct coverage if the injected signals were drawn directly from our priors, which is not the case here. Although **ms1**, **h4** and **s1y** are among the candidate EOSs upon which the GP was conditioned, they are not themselves synthetic EOSs drawn from the EOS prior process.




FIG. 5. Processes corresponding to the injection in Fig. 4 (event I with **h4** from Table I), showing the prior (cyan/blue) and posterior (magenta/red) for our model-agnostic (left) and model-informed (right) priors. Solid lines show the medians and shaded regions correspond to 50% and 90% credible regions for the one-dimensional marginal distributions for $p(\rho)$. Contours correspond to the 50% and 90% credible regions for the central density and pressure of M_1 and vertical black bars denote ρ_{nuc} , $2\rho_{\text{nuc}}$, and $6\rho_{\text{nuc}}$. We see stronger relative *a posteriori* constraints with the model-agnostic prior, whereas the posterior obtained with the model-informed prior is nearly indistinguishable from the prior itself.

Since the recovered posteriors for the injections depend on the GP prior, we may quantify the degree to which the GW data prefer one prior or the other via Bayesian model selection. We estimate the Bayes factor

$$B_{\mathcal{I}}^A = \frac{P(d|\{\varepsilon_{j^*}^{(\alpha)}\}, \vec{\sigma}_{\mathcal{A}}, \mathcal{H})}{P(d|\{\varepsilon_{j^*}^{(\alpha)}\}, \vec{\sigma}_{\mathcal{I}}, \mathcal{H})} \quad (32)$$

between our model-agnostic and model-informed priors for all events listed in Table I, calculating the evidences according to Eq. (31). For most simulated signals, neither prior is strongly preferred by the data, with the exception of injection I with **ms1**. As **ms1** is particularly stiff, it lies at the very edge of our model-informed prior's support, and therefore there is much more evidence in favor of the model-agnostic prior. This appears to be more prevalent for particularly stiff EOSs than for particularly soft ones (compare $B_{\mathcal{I}}^A$ for **ms1** and **sly**), but we typically find that the GW data do not disfavor our model-informed prior at high confidence. One could imagine performing a similar model selection scheme on different GP priors conditioned on specific features of the NS EOS, to test, e.g., for the existence of a first-order phase transition in the NS EOS.

VI. GW170817

Having tested our non-parametric inference on simulated GW events, we now apply it to real data from GW170817. We repeat the analysis of the previous section using the same model-agnostic and model-informed priors. Our recovered $M_{1,2}$ - $\Lambda_{1,2}$ posteriors and EOS posterior processes are shown in Figs. 6 and 7, respectively. One-dimensional credible regions for the component masses and tidal deformabilities are reported in Table I, which also presents the Bayes factor between the two priors.

With the model-agnostic prior, we find that GW170817 places significant *a posteriori* constraints on the tidal deformabilities, favoring a limited region corresponding to relatively soft EOSs. Indeed, we infer a chirp deformability of $\tilde{\Lambda} = 210_{-113}^{+383}$. Using $\Lambda(M; \varepsilon)$ for every synthetic EOS $\varepsilon^{(a)}(p)$, we find $\Lambda_{1,4} = 160_{-113}^{+448}$. The lower bound of our

Injected						Recovered with model-agnostic (model-informed) prior					B_T^A	
	$m_1 [M_\odot]$	$m_2 [M_\odot]$	EOS	Λ_1	Λ_2	$\tilde{\Lambda}$	$m_1 [M_\odot]$	$m_2 [M_\odot]$	Λ_1	Λ_2	$\tilde{\Lambda}$	
(I)	1.57	1.19	MS1	714	3320	1550	$1.40^{+0.27}_{-0.04}$	$1.30^{+0.05}_{-0.20}$	611^{+979}_{-276}	2350^{+2030}_{-897}	1530^{+387}_{-455}	22.6 ± 0.7
							$(1.72^{+0.03}_{-0.27})$	$(1.08^{+0.18}_{-0.03})$	(210^{+472}_{-70})	(3190^{+279}_{-2070})	(890^{+177}_{-112})	
			H4	390	2250	983	$1.57^{+0.12}_{-0.22}$	$1.15^{+0.18}_{-0.07}$	341^{+504}_{-265}	1200^{+1440}_{-488}	836^{+337}_{-366}	0.26 ± 0.01
							$(1.60^{+0.10}_{-0.23})$	$(1.14^{+0.17}_{-0.07})$	(231^{+413}_{-109})	(1152^{+1340}_{-407})	(740^{+167}_{-164})	
			SLy	130	834	354	$1.41^{+0.27}_{-0.05}$	$1.22^{+0.11}_{-0.13}$	155^{+363}_{-121}	792^{+842}_{-563}	463^{+277}_{-320}	0.44 ± 0.01
							$(1.40^{+0.24}_{-0.05})$	$(1.30^{+0.05}_{-0.18})$	(316^{+328}_{-192})	(862^{+1080}_{-281})	(667^{+147}_{-144})	
(II)	1.37	1.36	MS1	1580	1630	1600	$1.40^{+0.20}_{-0.05}$	$1.29^{+0.06}_{-0.15}$	636^{+637}_{-546}	1670^{+1330}_{-1028}	1200^{+541}_{-686}	0.85 ± 0.02
							$(1.40^{+0.23}_{-0.05})$	$(1.29^{+0.05}_{-0.17})$	(406^{+345}_{-226})	(1030^{+1210}_{-329})	(776^{+192}_{-199})	
			H4	1000	1034	1020	$1.39^{+0.22}_{-0.04}$	$1.31^{+0.04}_{-0.17}$	341^{+527}_{-257}	942^{+1240}_{-452}	721^{+461}_{-386}	0.30 ± 0.01
							$(1.39^{+0.23}_{-0.04})$	$(1.31^{+0.04}_{-0.18})$	(306^{+440}_{-121})	(942^{+1080}_{-276})	(728^{+178}_{-162})	
			SLy	349	362	355	$1.40^{+0.17}_{-0.05}$	$1.29^{+0.06}_{-0.13}$	140^{+278}_{-109}	341^{+581}_{-217}	247^{+332}_{-162}	1.40 ± 0.03
							$(1.41^{+0.13}_{-0.06})$	$(1.29^{+0.06}_{-0.10})$	(431^{+210}_{-202})	(902^{+547}_{-304})	(655^{+161}_{-130})	
(III)	1.42	1.33	MS1	1270	1850	1530	$1.62^{+0.13}_{-0.20}$	$1.14^{+0.17}_{-0.07}$	361^{+725}_{-278}	2550^{+1620}_{-1470}	1210^{+458}_{-474}	1.97 ± 0.06
							$(1.67^{+0.09}_{-0.19})$	$(1.13^{+0.13}_{-0.06})$	(246^{+194}_{-121})	(1830^{+1200}_{-895})	(782^{+191}_{-229})	
			H4	782	1190	968	$1.58^{+0.16}_{-0.16}$	$1.14^{+0.16}_{-0.07}$	251^{+605}_{-193}	2010^{+1409}_{-1130}	998^{+376}_{-515}	0.44 ± 0.01
							$(1.63^{+0.12}_{-0.17})$	$(1.13^{+0.14}_{-0.06})$	(210^{+308}_{-94})	(1960^{+713}_{-1050})	(697^{+190}_{-154})	
			SLy	269	420	338	$1.62^{+0.10}_{-0.24}$	$1.15^{+0.17}_{-0.08}$	110^{+333}_{-92}	641^{+1020}_{-480}	337^{+359}_{-230}	0.53 ± 0.01
							$(1.46^{+0.25}_{-0.09})$	$(1.17^{+0.16}_{-0.09})$	(231^{+331}_{-121})	(1040^{+1030}_{-458})	(637^{+133}_{-139})	
GW170817							$1.46^{+0.16}_{-0.09}$	$1.26^{+0.10}_{-0.10}$	110^{+278}_{-81}	361^{+529}_{-252}	210^{+383}_{-113}	1.12 ± 0.06
							$(1.41^{+0.15}_{-0.04})$	$(1.32^{+0.04}_{-0.13})$	(426^{+214}_{-229})	(772^{+641}_{-212})	(631^{+164}_{-122})	

TABLE I. Injected parameters used to test EOS recovery and the associated one-dimensional posterior credible regions. Three mass ratios and three injected EOSs are considered. We report the 90% credible regions for the model-agnostic (model-informed) prior along with the maximum *a posteriori* for each parameter separately. Additionally, we report the Bayes factor between our priors (B_T^A), generally showing no strong preference for either prior with a single event. Error estimates correspond to 1- σ uncertainty for Bayes factors. Analogous results for GW170817 are shown as well, with similar conclusions.

confidence interval suggests some of the smallest tidal deformabilities of any study to date. Our findings agree with the results of previous analyses, namely $\tilde{\Lambda} = 300^{+420}_{-230}$ [28] (highest-posterior-density 90% confidence interval about the median), $\tilde{\Lambda} = 222^{+420}_{-138}$ [14] and $\Lambda_{1.4} = 190^{+390}_{-120}$ [15] (both symmetric 90% confidence intervals about the median).⁶ Moreover, the 90% confidence region of our Λ_1 - Λ_2 posterior, shown in Fig. 6, is markedly similar to the corresponding region of the posterior obtained with the spectral method (cf. Fig. 1 of Ref. [15]).

From Fig. 7, we observe a preference for a slightly softer-than-average pressure-density relation at mid-range densities. In particular, we obtain the 90%-credible regions $p(2\rho_{\text{nuc}}) = 1.35^{+1.8}_{-1.2} \times 10^{34}$ dyn/cm² and $p(6\rho_{\text{nuc}}) = 8.86^{+4.3}_{-5.9} \times 10^{35}$ dyn/cm² on the pressure at twice and six times nuclear density. These values agree with the median pressures $p(2\rho_{\text{nuc}}) = 3.5^{+2.7}_{-1.7} \times 10^{34}$ dyn/cm² and $p(6\rho_{\text{nuc}}) = 9.0^{+7.9}_{-2.6} \times 10^{35}$ dyn/cm² (90% confidence) obtained in Ref. [15]. Furthermore, Fig. 7 displays our inference for the central density $\rho_{c,1}$ and pressure $p_{c,1}$ of the more massive component of the binary; we find that the posterior peaks at $\rho_{c,1} \approx 3\rho_{\text{nuc}}$. Since $M(\rho_c; \varepsilon)$ is known for each synthetic EOS, the posterior process can also be used to make an inference for the maximum mass supported by the NS EOS. We find that $M_{\text{max}} = 2.09^{+0.37}_{-0.16} M_\odot$ with the model-agnostic prior.

Repeating the inference with the model-informed prior, we recover a smaller *a posteriori* 90% confidence region that has less support for soft EOSs, as can be seen in the right-hand panel of Fig. 6. Our inferred tidal deformabilities are thus somewhat larger: $\tilde{\Lambda} = 631^{+164}_{-122}$ and $\Lambda_{1.4} = 556^{+163}_{-172}$. The upper bounds on the tidal deformabilities are similar to the results of the original analyses of GW170817 [13, 28], but the lower bounds are much larger. This is because the tabulated EOSs used to construct the model-informed prior cannot produce $\Lambda_{1.4} \lesssim 290$ (the `slly` value), while the GW data appear to support deformabilities $\Lambda_{1.4} \lesssim 200$. Electromagnetic observations of the kilonova associated with GW170817 have found similarly large lower bounds for $\tilde{\Lambda}$ [51, 55, 81]. The numerical relativity simulations of Ref. [81], in particular, suggest $\tilde{\Lambda} \gtrsim 300$, in mild tension with our model-agnostic result but more compatible with our

⁶ Because the results reported in the literature are generally symmetric confidence intervals about the median, whereas ours are highest-posterior-density intervals about the maximum *a posteriori*, our confidence intervals are systematically shifted downwards relative to them (see, for example, Fig. 11 of Ref. [28]).

model-informed analysis. Curiously, the $\Lambda_{1.4}$ 90% confidence region for the model-agnostic prior includes a long tail to large values and contains the maximum *a posteriori* estimate from the model-informed prior, although the 50% confidence region is much tighter.

The EOS posterior process for the model-informed prior is plotted in the right-hand panel of Fig. 7. The pressure constraints extracted from the posterior are $p(2\rho_{\text{nuc}}) = 4.73_{-2.5}^{+1.4} \times 10^{34}$ dyn/cm² and $p(6\rho_{\text{nuc}}) = 7.55_{-3.2}^{+2.0} \times 10^{35}$ dyn/cm², consistent with Ref. [15] (see Fig. 3). The posterior for the central density $\rho_{c,1}$ peaks around $2\rho_{\text{nuc}}$, and the constraint on the maximum mass supported is tightened to $M_{\text{max}} = 2.04_{-0.002}^{+0.22} M_{\odot}$.

Overall, our results for GW170817 are consistent with the literature [13–15, 28], favoring relatively low $\tilde{\Lambda}$ and a correspondingly soft EOS. In fact, Fig. 3 shows that our inferred model-agnostic posterior process for $\varepsilon(p)$ favors a slightly softer EOS than that found with the spectral method [15], though the uncertainties are quite broad. On the other hand, the model-informed posterior prefers slightly larger pressures than the spectral method’s. Nonetheless, compared to our priors—i.e. $p_{\text{prior}}(2\rho_{\text{nuc}}) = 1.09_{-1.1}^{+8.8} \times 10^{34}$ ($4.96_{-2.9}^{+1.3} \times 10^{34}$) dyn/cm² for the model-agnostic (model-informed) prior—we generally see smaller 90% confidence regions and shifted modes in the pressure-density plane *a posteriori*.

Although the differences between prior and posterior are more pronounced in the model-agnostic case, neither prior is strongly favored by the data. We find $B_{\text{T}}^{\text{A}} = 1.12 \pm 0.06$ (point estimate and 1- σ uncertainty), suggesting that our inference of the tidal deformabilities and the EOS itself may be relatively prior-dominated. In other words, the data from GW170817 alone cannot strongly differentiate between a model that yields $p(2\rho_{\text{nuc}}) = 1.35_{-1.2}^{+1.8} \times 10^{34}$ dyn/cm² and one that yields $p(2\rho_{\text{nuc}}) = 4.73_{-2.5}^{+1.4} \times 10^{34}$ dyn/cm², suggesting that all such constraints should be interpreted with care. We expect that combined data from multiple GW signals will be needed to discriminate between different models for the EOS.

VII. CONCLUSIONS

GW measurements of the macroscopic properties of NSs offer a promising means of deducing information about the nuclear microphysics that governs their internal structure. The non-parametric framework for NS EOS inference developed here provides direct constraints on the pressure-density relation inside the star from the standard GW likelihood. The GP formulation of the EOS can compactly and faithfully represent a more diverse set of possible EOSs than parametric models, and since its systematic uncertainties are independent of those from previous analyses, non-parametric inference can be used to complement and corroborate existing approaches. The full EOS posterior process produced by our approach allows pressures and densities within the NS to be inferred, derived properties like the maximum NS mass to be constrained, and model selection between various EOS priors to be performed.

We studied simulated GW170817-like signals using real detector noise as a demonstration of our method’s efficacy at recovering a known EOS. The inference produced informative constraints on the tidal deformabilities and the pressure-density relation, in the sense that the 90% confidence region of the posterior was reduced. Overall, we find that injections with stiffer EOSs favor stiffer EOSs *a posteriori*, and softer injected EOSs favor softer posteriors, as expected. However, the constraints depend somewhat strongly on the prior assumed, and the data typically do not favor the model-agnostic prior over the model-informed one, or vice-versa.

Using the same model-informed and model-agnostic priors, we analyzed GW170817 to infer tidal deformabilities, pressures and the maximum mass supported by the EOS. We summarize the results below and discuss their significance.

Tidal deformabilities. Analyzing GW170817 with the model-agnostic prior, we infer a canonical deformability of $\Lambda_{1.4} = 160_{-113}^{+448}$ and a chirp deformability of $\tilde{\Lambda} = 210_{-113}^{+383}$, in comfortable agreement with previous analyses [14, 15]. The maximum *a posteriori* values we obtain for the canonical and chirp deformabilities are the smallest claimed to date, but come with broad uncertainties. Such small tidal deformabilities favor an EOS that resembles the softest nuclear-theoretic models available. Repeating the analysis with the model-informed prior, the inferred tidal deformabilities are $\Lambda_{1.4} = 556_{-172}^{+163}$ and $\tilde{\Lambda} = 631_{-122}^{+164}$. The 90% confidence regions from both priors overlap, although the maxima *a posteriori* are notably different.

Pressures. The posterior process places direct constraints on the EOS in the pressure-density plane, enabling us to extract the constraints $p(2\rho_{\text{nuc}}) = 1.35_{-1.2}^{+1.8} \times 10^{34}$ dyn/cm² and $p(6\rho_{\text{nuc}}) = 8.86_{-5.9}^{+4.3} \times 10^{35}$ dyn/cm² on the pressure at twice and six times the nuclear saturation density with the model-agnostic prior. These figures are consistent with the pressure constraints calculated with the spectral method. Qualitatively, we observe a similar preference for lower-than-average pressures at mid-range densities as was noted in Ref. [15]. The pressures inferred with the model-informed prior are $p(2\rho_{\text{nuc}}) = 4.73_{-2.5}^{+1.4} \times 10^{34}$ dyn/cm² and $p(6\rho_{\text{nuc}}) = 7.55_{-3.2}^{+2.0} \times 10^{35}$ dyn/cm², favoring a slightly stiffer EOS. Nonetheless, the results of the inference are still consistent with those of previous analyses [15] within their uncertainties.

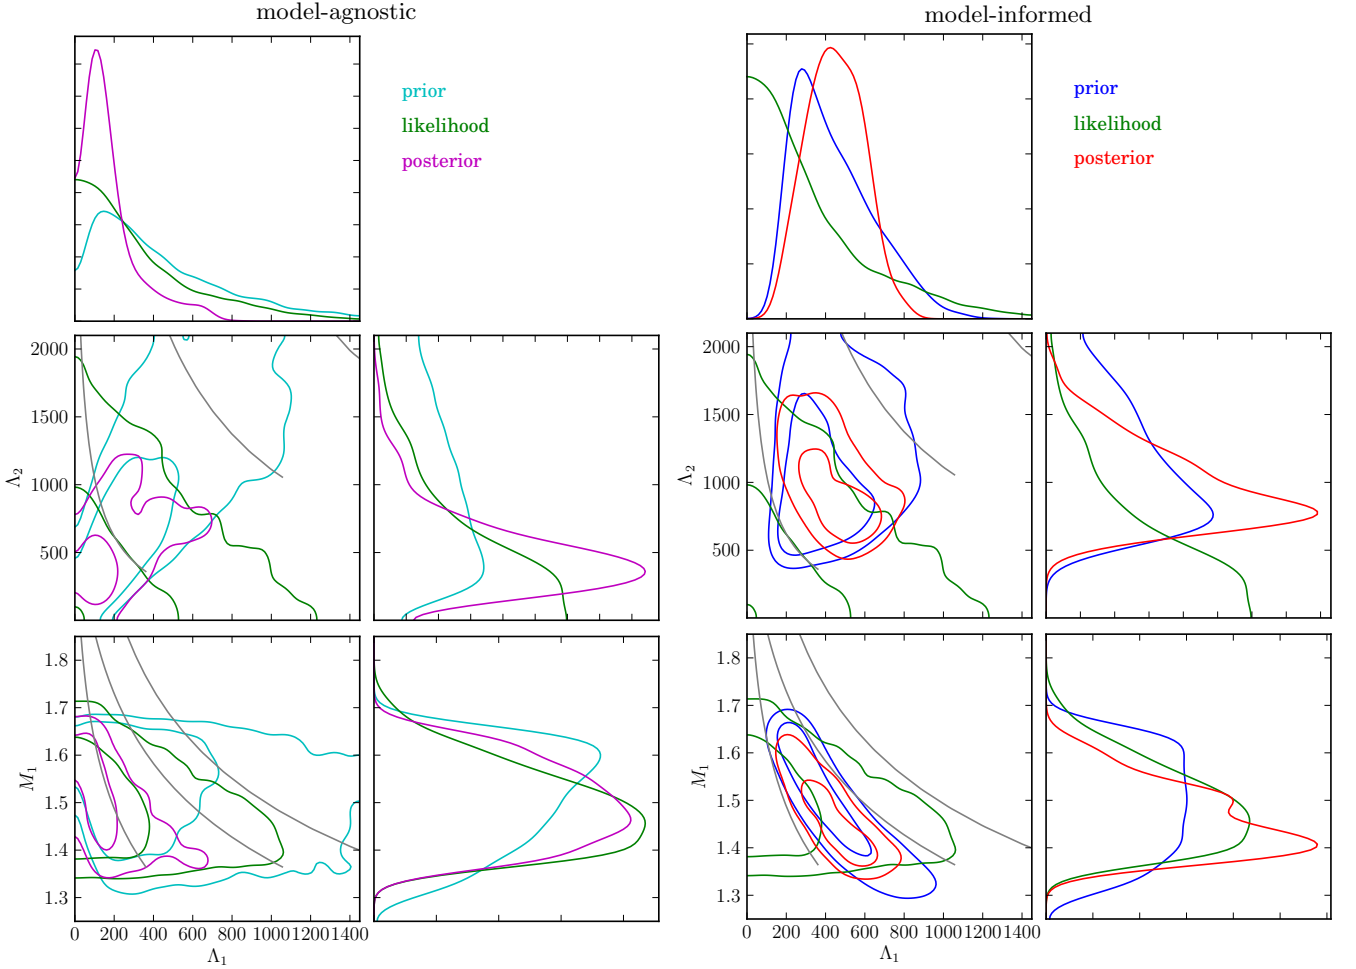


FIG. 6. Marginal distributions for GW170817, analogous to Fig. 4. The priors (cyan/blue), likelihood (green), and posteriors (magenta/red) are shown for both the model-agnostic (left) and model-informed (right) priors. Reference curves (gray) are shown for (from left to right) **s1y**, **h4**, and **ms1**. Contours in the joint distributions correspond to the 50% and 90% credible regions, and we see that the model-agnostic prior’s *a posteriori* constraints are similar to those reported in Ref. [15], whereas the model-informed prior favors significantly larger $\Lambda_{1,2}$ but similar masses.

Maximum mass. The EOS posterior process gives us access to maximum-mass information through correlations between the high densities relevant for M_{max} and the mid-range densities probed by GW170817. With the model-agnostic prior, we constrain the maximum NS mass to be $M_{\text{max}} = 2.09^{+0.37}_{-0.16} M_{\odot}$, disfavoring especially stiff candidate EOSs that produce $M_{\text{max}} \gtrsim 2.5 M_{\odot}$. Ours is the first reported constraint derived exclusively from GW data. This is compatible with multimessenger studies of GW170817 via EOS model selection [82] and approximate universal relations [83], numerical relativity [84], and general-relativistic magnetohydrodynamic simulations [85], which account for the observed spectrum of the electromagnetic counterpart, the lifetime of the merger remnant, and the GW-inferred source properties. Taken together, the multimessenger studies suggest $2.15 M_{\odot} \lesssim M_{\text{max}} \lesssim 2.28 M_{\odot}$. The constraint on the maximum NS mass is tightened to $M_{\text{max}} = 2.04^{+0.22}_{-0.002} M_{\odot}$ with the model-informed prior.

Bayes factor. Despite producing noticeably dissimilar *a posteriori* credible regions for the tidal deformabilities and the pressure-density relation, the Bayes factors between the model-agnostic and model-informed priors are $O(1)$ for virtually all of the simulated events considered. For GW170817, we find $B_{\mathcal{I}}^A = 1.12 \pm 0.06$ (point estimate and 1- σ uncertainty). This implies that data from a single event are not sufficiently informative to strongly favor either model. We conclude that EOS constraints derived from GW170817 alone are likely to be sensitive to the choice of prior. Thus, the EOS priors must be carefully motivated on physical grounds and associated constraints interpreted scrupulously.

Given the prior-dependence of the results, revisiting the choice of candidate EOSs used to condition the GP priors may be worthwhile. One could expand the training set to include candidate EOSs softer than **s1y** and stiffer than **ms1**. Since early indications are that the NS EOS is relatively soft, extending the prior at its soft edge is particularly

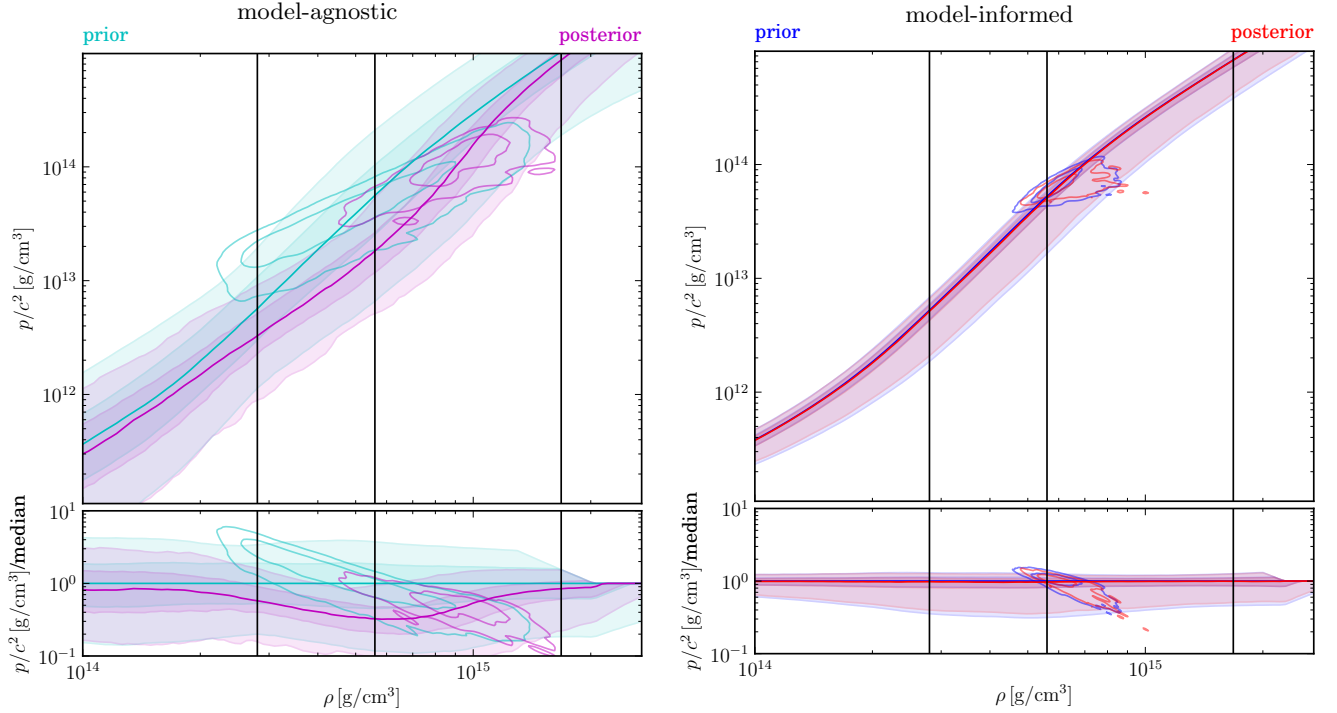


FIG. 7. Processes for the model-agnostic (*left*) and model-informed (*right*) priors with GW170817, both *a posteriori* (magenta/red) and *a priori* (cyan/blue). We find only marginally tightened posterior constraints with the model-informed prior, whereas the model-agnostic posterior clearly favors softer EOSs. As in Fig. 5, solid lines show medians, shaded regions represent 50% and 90% credible regions for $p(\rho)$, contours represent the 50% and 90% credible regions for the central densities and pressures of M_1 , and vertical bars denote ρ_{nuc} , $2\rho_{\text{nuc}}$, and $6\rho_{\text{nuc}}$.

important for obtaining reliable model-informed inferences from GW data. Adding candidate EOSs with, e.g., first order phase transitions or strange quark matter could also make the training set more fully representative of the theoretical possibilities for the EOS. Similarly, one could broaden the training set to include nonbarotropic candidate EOSs, such as recent proposals that depend on the environmental dark matter density [86–88]. Despite already covering a large area of the pressure-density plane, the model-agnostic prior could also benefit from conditioning on an expanded and more representative set of input EOSs, as it would assimilate an even greater degree of variation in the functional behavior of the EOS. Coupled with improved hyperparameter selection via automated marginalization over many possible settings (cf. Eq. (11)), it may be possible to deliver tighter constraints on the tidal deformabilities and the pressure-density relation with priors better grounded in nuclear theory.

Although we found that posterior constraints from a single GW event are rather broad and prior-dependent, we expect that the impact of prior assumptions will eventually be overcome by the combined data from many GW signals. Our Monte Carlo integration scheme can be straightforwardly extended to incorporate an arbitrary number of events: we simply draw masses (e.g., $M_{1,A}^{(a)}$) consistent with each event A separately for each synthetic EOS, and then combine the likelihoods \mathcal{L}_A from each event so that the overall weight for the a^{th} synthetic EOS is $\mathcal{L}^{(a)} = \prod_A \mathcal{L}_A(d_A | M_{1,A}^{(a)}, M_{2,A}^{(a)}, \Lambda_{1,A}^{(a)}, \Lambda_{2,A}^{(a)}; \mathcal{H}_A)$. Electromagnetic observations of NS radii from NICER and moment of inertia measurements from radio observations can also be incorporated as additional likelihoods when available. A precise quantification of the non-parametric method’s ability to constrain the EOS with multiple GW signals will be investigated in future work.

As illustrated by our Bayes factors between the model-agnostic and model-informed priors, our method naturally provides evidence estimates which can be used for Bayesian hypothesis ranking. Leveraging the flexibility that the non-parametric representation offers when selecting the EOS prior, one could perform diverse kinds of model selection beyond the simple example presented here. For instance, evidence for different NS matter compositions (e.g. purely hadronic vs. hyperonic) or phenomenological features of the EOS (e.g. strong phase transitions) could be compared. This can be done by changing the set of tabulated EOSs used to condition the underlying GP, so as to construct separate priors reflecting the typical behavior of different classes of EOS. In the context of GW170817, a hypothesis ranking scheme could also be applied to determine whether the GW data identify the event as a binary NS, a binary black hole, or a NS-black hole merger. Model selection with a GP representation for the EOS has the potential to

shed light on a variety of interesting questions about the macroscopic properties and internal structure of NSs. More broadly, the versatility of non-parametric EOS inference will help maximize the information obtained from future observations, serving as a key component in the effort to constrain the supranuclear EOS.

ACKNOWLEDGMENTS

The authors thank Daniel Holz, Shaon Ghosh and the members of the LVC Extreme Matter group for useful discussions about this work. Computational resources provided by the LIGO Laboratory and supported by National Science Foundation Grants PHY-0757058 and PHY-0823459 are gratefully acknowledged. P. L. is supported in part by the Natural Sciences and Engineering Research Council of Canada, and by NSF grants PHY 15-05124 and PHY 17-08081 to the University of Chicago. R. E. is supported at the University of Chicago by the Kavli Institute for Cosmological Physics through an endowment from the Kavli Foundation and its founder Fred Kavli.

-
- [1] J. L. Friedman and N. Stergioulas, *Rotating Relativistic Stars* (Cambridge University Press, 2013).
 - [2] J. M. Lattimer, “The Nuclear Equation of State and Neutron Star Masses,” *Annu. Rev. of Nucl. Part. Sci.* **62**, 485–515 (2012), arXiv:1305.3510 [nucl-th].
 - [3] L. Lindblom, “Determining the nuclear equation of state from neutron-star masses and radii,” *Astrophys. J.* **398**, 569–573 (1992).
 - [4] J. M. Lattimer and M. Prakash, “Neutron Star Structure and the Equation of State,” *Astrophys. J.* **550**, 426–442 (2001), astro-ph/0002232.
 - [5] F. Özel and P. Freire, “Masses, Radii, and the Equation of State of Neutron Stars,” *Annu. Rev. Astron. Astrophys.* **54**, 401–440 (2016), arXiv:1603.02698 [astro-ph.HE].
 - [6] I. A. Morrison, T. W. Baumgarte, S. L. Shapiro, and V. R. Pandharipande, “The Moment of Inertia of the Binary Pulsar J0737-3039A: Constraining the Nuclear Equation of State,” *Astrophys. J. Lett.* **617**, L135–L138 (2004), astro-ph/0411353.
 - [7] J. M. Lattimer and B. F. Schutz, “Constraining the Equation of State with Moment of Inertia Measurements,” *Astrophys. J.* **629**, 979–984 (2005), astro-ph/0411470.
 - [8] M. Kramer and N. Wex, “TOPICAL REVIEW: The double pulsar system: a unique laboratory for gravity,” *Class. Quantum Grav.* **26**, 073001 (2009).
 - [9] P. E. Dewdney, P. J. Hall, R. T. Schilizzi, and T. J. L. W. Lazio, “The Square Kilometre Array,” *IEEE Proceedings* **97**, 1482–1496 (2009).
 - [10] K. C. Gendreau, Z. Arzoumanian, and T. Okajima, “The Neutron star Interior Composition ExploreR (NICER): an Explorer mission of opportunity for soft x-ray timing spectroscopy,” in *Space Telescopes and Instrumentation 2012: Ultraviolet to Gamma Ray*, Proc. SPIE, Vol. 8443 (2012) p. 844313.
 - [11] J. Aasi *et al.* (LIGO Scientific Collaboration), “Advanced LIGO,” *Class. Quantum Grav.* **32**, 074001 (2015), arXiv:1411.4547 [gr-qc].
 - [12] F. Acernese *et al.* (Virgo Collaboration), “Advanced Virgo: a second-generation interferometric gravitational wave detector,” *Class. Quantum Grav.* **32**, 024001 (2015), arXiv:1408.3978 [gr-qc].
 - [13] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), “GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral,” *Phys. Rev. Lett.* **119**, 161101 (2017), arXiv:1710.05832 [gr-qc].
 - [14] Soumi De *et al.*, “Tidal Deformabilities and Radii of Neutron Stars from the Observation of GW170817,” *Phys. Rev. Lett.* **121**, 091102 (2018), arXiv:1804.08583 [astro-ph.HE].
 - [15] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), “GW170817: Measurements of neutron star radii and equation of state,” *ArXiv e-prints*, arXiv:1805.11581 (2018), arXiv:1805.11581 [gr-qc].
 - [16] C. S. Kochanek, “Coalescing binary neutron stars,” *Astrophys. J.* **398**, 234–247 (1992).
 - [17] D. Lai and A. G. Wiseman, “Innermost stable circular orbit of inspiraling neutron-star binaries: Tidal effects, post-Newtonian effects, and the neutron-star equation of state,” *Phys. Rev. D* **54**, 3958–3964 (1996), gr-qc/9609014.
 - [18] É. É. Flanagan and T. Hinderer, “Constraining neutron-star tidal Love numbers with gravitational-wave detectors,” *Phys. Rev. D* **77**, 021502 (2008), arXiv:0709.1915.
 - [19] T. Hinderer, B. D. Lackey, R. N. Lang, and J. S. Read, “Tidal deformability of neutron stars with realistic equations of state and their gravitational wave signatures in binary inspiral,” *Phys. Rev. D* **81**, 123016 (2010), arXiv:0911.3535 [astro-ph.HE].
 - [20] J. S. Read *et al.*, “Matter effects on binary neutron star waveforms,” *Phys. Rev. D* **88**, 044042 (2013), arXiv:1306.4065 [gr-qc].
 - [21] W. Del Pozzo, T. G. F. Li, M. Agathos, C. Van Den Broeck, and S. Vitale, “Demonstrating the Feasibility of Probing the Neutron-Star Equation of State with Second-Generation Gravitational-Wave Detectors,” *Phys. Rev. Lett.* **111**, 071101 (2013), arXiv:1307.8338 [gr-qc].
 - [22] L. Wade *et al.*, “Systematic and statistical errors in a Bayesian approach to the estimation of the neutron-star equation of state using advanced gravitational wave detectors,” *Phys. Rev. D* **89**, 103012 (2014), arXiv:1402.5156 [gr-qc].

- [23] T. Damour, A. Nagar, and L. Villain, “Measurability of the tidal polarizability of neutron stars in late-inspiral gravitational-wave signals,” *Phys. Rev. D* **85**, 123007 (2012), arXiv:1203.4352 [gr-qc].
- [24] J. Veitch *et al.*, “Parameter estimation for compact binaries with ground-based gravitational-wave observations using the LALInference software library,” *Phys. Rev. D* **91**, 042003 (2015), arXiv:1409.7215 [gr-qc].
- [25] M. Favata, “Systematic Parameter Errors in Inspiring Neutron Star Binaries,” *Phys. Rev. Lett.* **112**, 101101 (2014), arXiv:1310.8288 [gr-qc].
- [26] M. Hannam, D. A. Brown, S. Fairhurst, C. L. Fryer, and I. W. Harry, “When can Gravitational-wave Observations Distinguish between Black Holes and Neutron Stars?” *Astrophys. J. Lett.* **766**, L14 (2013), arXiv:1301.5616 [gr-qc].
- [27] P. Landry and B. Kumar, “Constraints on the Moment of Inertia of PSR J0737-3039A from GW170817,” *Astrophys. J. Lett.* **868**, L22 (2018), arXiv:1807.04727 [gr-qc].
- [28] Abbott, B. P. *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), “Properties of the binary neutron star merger GW170817,” *ArXiv e-prints* (2018), arXiv:1805.11579 [gr-qc].
- [29] F. Douchin and P. Haensel, “A unified equation of state of dense matter and neutron star structure,” *Astron. Astrophys.* **380**, 151–167 (2001), astro-ph/0111092.
- [30] H. Müller and B. D. Serot, “Relativistic mean-field theory and the high-density nuclear equation of state,” *Nuclear Physics A* **606**, 508–537 (1996), nucl-th/9603037.
- [31] K. Yagi and N. Yunes, “Binary Love relations,” *Class. Quantum Grav.* **33**, 13LT01 (2016), arXiv:1512.02639 [gr-qc].
- [32] K. Yagi and N. Yunes, “Approximate universal relations among tidal parameters for neutron star binaries,” *Class. Quantum Grav.* **34**, 015006 (2017), arXiv:1608.06187 [gr-qc].
- [33] K. Chatziioannou, C.-J. Haster, and A. Zimmerman, “Measuring the neutron star tidal deformability with equation-of-state-independent relations and gravitational waves,” *Phys. Rev. D* **97**, 104036 (2018), arXiv:1804.03221 [gr-qc].
- [34] L. Lindblom, “Causal representations of neutron-star equations of state,” *Phys. Rev. D* **97**, 123019 (2018), arXiv:1804.04072 [astro-ph.HE].
- [35] L. Lindblom, “Spectral representations of neutron-star equations of state,” *Phys. Rev. D* **82**, 103011 (2010), arXiv:1009.0738 [astro-ph.HE].
- [36] L. Lindblom and N. M. Indik, “Spectral approach to the relativistic inverse stellar structure problem,” *Phys. Rev. D* **86**, 084003 (2012), arXiv:1207.3744 [astro-ph.HE].
- [37] L. Lindblom and N. M. Indik, “Spectral approach to the relativistic inverse stellar structure problem II,” *Phys. Rev. D* **89**, 064003 (2014), arXiv:1310.0803 [astro-ph.HE].
- [38] J. S. Read, B. D. Lackey, B. J. Owen, and J. L. Friedman, “Constraints on a phenomenologically parametrized neutron-star equation of state,” *Phys. Rev. D* **79**, 124032 (2009), arXiv:0812.2163.
- [39] B. D. Lackey and L. Wade, “Reconstructing the neutron-star equation of state with gravitational-wave detectors from a realistic population of inspiralling binary neutron stars,” *Phys. Rev. D* **91**, 043002 (2015), arXiv:1410.8866 [gr-qc].
- [40] G. Raaijmakers, T. E. Riley, and A. L. Watts, “A pitfall of piecewise-polytropic equation of state inference,” *Mon. Not. Royal Astron.* **478**, 2177–2192 (2018), arXiv:1804.09087 [astro-ph.HE].
- [41] M. F. Carney, L. E. Wade, and B. S. Irwin, “Comparing two models for measuring the neutron star equation of state from gravitational-wave signals,” *Phys. Rev. D* **98**, 063004 (2018), arXiv:1805.11217 [gr-qc].
- [42] M. Agathos *et al.*, “Constraining the neutron star equation of state with gravitational wave signals from coalescing binary neutron stars,” *Phys. Rev. D* **92**, 023012 (2015), arXiv:1503.05405 [gr-qc].
- [43] T. Abdelsalhin, A. Maselli, and V. Ferrari, “Solving the relativistic inverse stellar problem through gravitational waves observation of binary neutron stars,” *ArXiv e-prints* (2017), arXiv:1712.01303 [gr-qc].
- [44] S. Han and A. W. Steiner, “Tidal deformability with sharp phase transitions in (binary) neutron stars,” *ArXiv e-prints* (2018), arXiv:1810.10967 [nucl-th].
- [45] M. Alford, M. Braby, M. Paris, and S. Reddy, “Hybrid Stars that Masquerade as Neutron Stars,” *Astrophys. J.* **629**, 969–978 (2005), nucl-th/0411016.
- [46] T. Miyatsu, T. Kambe, and K. Saito, “Equation of state for hybrid stars with strangeness,” *ArXiv e-prints* (2017), arXiv:1702.05871 [nucl-th].
- [47] T. Fischer *et al.*, “Quark deconfinement as a supernova explosion engine for massive blue supergiant stars,” *Nature Astronomy* (2018), 10.1038/s41550-018-0583-0, arXiv:1712.08788 [astro-ph.HE].
- [48] P. Haensel, A. Y. Potekhin, and D. G. Yakovlev, *Neutron Stars 1 : Equation of State and Structure*, Astrophysics and Space Science Library, Vol. 326 (Springer-Verlag, 2007).
- [49] LIGO Scientific Collaboration, “LIGO Algorithm Library - LALSuite,” free software (GPL) (2018).
- [50] E. R. Most, L. R. Weih, L. Rezzolla, and J. Schaffner-Bielich, “New Constraints on Radii and Tidal Deformabilities of Neutron Stars from GW170817,” *Physical Review Letters* **120**, 261103 (2018), arXiv:1803.00549 [gr-qc].
- [51] D. Radice, A. Perego, F. Zappa, and S. Bernuzzi, “GW170817: Joint Constraint on the Neutron Star Equation of State from Multimessenger Observations,” *Astrophys. J. Lett.* **852**, L29 (2018), arXiv:1711.03647 [astro-ph.HE].
- [52] B. P. Abbott *et al.*, “Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA,” *Living Reviews in Relativity* **21**, 3 (2018).
- [53] Z. Doctor, B. Farr, D. E. Holz, and M. Pürrer, “Statistical gravitational waveform models: What to simulate next?” *Phys. Rev. D* **96**, 123011 (2017), arXiv:1706.05408 [astro-ph.HE].
- [54] E. A. Huerta *et al.*, “Eccentric, nonspinning, inspiral, Gaussian-process merger approximant for the detection and characterization of eccentric binary black hole mergers,” *Phys. Rev. D* **97**, 024031 (2018), arXiv:1711.06276 [gr-qc].
- [55] M. W. Coughlin *et al.*, “Constraints on the neutron star equation of state from AT2017gfo using radiative transfer simulations,” *Mon. Not. Royal Astron.* **480**, 3871–3878 (2018), arXiv:1805.09371 [astro-ph.HE].

- [56] C. J. Moore and J. R. Gair, “Novel Method for Incorporating Model Uncertainties into Gravitational Wave Parameter Estimates,” *Physical Review Letters* **113**, 251101 (2014), arXiv:1412.3657 [gr-qc].
- [57] C. J. Moore, C. P. L. Berry, A. J. K. Chua, and J. R. Gair, “Improving gravitational-wave parameter estimation using Gaussian process regression,” *Phys. Rev. D* **93**, 064001 (2016), arXiv:1509.04066 [gr-qc].
- [58] J. Lange, R. O’Shaughnessy, and M. Rizzo, “Rapid and accurate parameter inference for coalescing, precessing compact binaries,” *ArXiv e-prints* (2018), arXiv:1805.10457 [gr-qc].
- [59] S. R. Taylor and D. Gerosa, “Mining gravitational-wave catalogs to understand binary stellar evolution: A new hierarchical Bayesian framework,” *Phys. Rev. D* **98**, 083017 (2018), arXiv:1806.08365 [astro-ph.HE].
- [60] S. Vitale and W. M. Farr, “Measuring the star formation rate with gravitational waves from binary black holes,” *ArXiv e-prints* (2018), arXiv:1808.00901 [astro-ph.HE].
- [61] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning* (The MIT Press, 2006).
- [62] D. Nguyen-Tuong, J. Peters, and M. Seeger, “Local Gaussian Process Regression for Real Time Online Model Learning and Control,” in *Proceedings of the 21st International Conference on Neural Information Processing Systems, NIPS’08* (Curran Associates Inc., USA, 2008) pp. 1193–1200.
- [63] F. Meier, P. Hennig, and S. Schaal, “Local Gaussian Regression,” *CoRR* **abs/1402.0645** (2014), arXiv:1402.0645.
- [64] R. C. Tolman, “Static Solutions of Einstein’s Field Equations for Spheres of Fluid,” *Physical Review* **55**, 364–373 (1939).
- [65] J. R. Oppenheimer and G. M. Volkoff, “On Massive Neutron Cores,” *Physical Review* **55**, 374–381 (1939).
- [66] P. Landry and E. Poisson, “Relativistic theory of surficial Love numbers,” *Phys. Rev. D* **89**, 124011 (2014), arXiv:1404.6798 [gr-qc].
- [67] B. K. Harrison, K. S. Thorne, M. Wakano, and J. A. Wheeler, *Gravitation Theory and Gravitational Collapse* (University of Chicago Press, 1965).
- [68] LIGO Scientific Collaboration and Virgo Collaboration, “Properties of the binary neutron star merger GW170817,” <https://dcc.ligo.org/LIGO-P1800061/public/> (2018).
- [69] C. Wang and R. M. Neal, “MCMC methods for Gaussian process models using fast approximations for the likelihood,” *ArXiv e-prints* (2013), arXiv:1305.2235 [stat.CO].
- [70] M. K. Titsias, M. Rattray, and N. D. Lawrence, “Markov chain Monte Carlo algorithms for Gaussian processes,” in *Bayesian Time Series Models*, edited by D. Barber, A. T. Cemgil, and S. Chiappa (Cambridge University Press, 2011) p. 295316.
- [71] Xtreme Astrophysics Group, “Dense Matter EoS,” <http://xtreme.as.arizona.edu/NeutronStars/>.
- [72] M. Alford, M. Braby, M. Paris, and S. Reddy, “Hybrid Stars that Masquerade as Neutron Stars,” *Astrophys. J.* **629**, 969–978 (2005), nucl-th/0411016.
- [73] L. Engvik, E. Osnes, M. Hjorth-Jensen, G. Bao, and E. Ostgaard, “Asymmetric Nuclear Matter and Neutron Star Properties,” *Astrophys. J.* **469**, 794 (1996), nucl-th/9509016.
- [74] B. D. Lackey, M. Nayyar, and B. J. Owen, “Observational constraints on hyperons in neutron stars,” *Phys. Rev. D* **73**, 024021 (2006), astro-ph/0507312.
- [75] H. M  ther, M. Prakash, and T. L. Ainsworth, “The nuclear symmetry energy in relativistic Brueckner-Hartree-Fock calculations,” *Phys. Lett. B* **199**, 469–474 (1987).
- [76] H. M  ller and B. D. Serot, “Relativistic mean-field theory and the high-density nuclear equation of state,” *Nucl. Phys. A* **606**, 508–537 (1996), nucl-th/9603037.
- [77] F. Douchin and P. Haensel, “A unified equation of state of dense matter and neutron star structure,” *Astron. Astrophys.* **380**, 151–167 (2001), astro-ph/0111092.
- [78] J. Antoniadis *et al.*, “A Massive Pulsar in a Compact Relativistic Binary,” *Science* **340**, 448 (2013), arXiv:1304.6875 [astro-ph.HE].
- [79] LIGO Scientific Collaboration and Virgo Collaboration, “Data release for event GW170817,” <https://www.gw-openscience.org/events/GW170817/> (2017).
- [80] A. Buonanno, B. R. Iyer, E. Ochsner, Y. Pan, and B. S. Sathyaprakash, “Comparison of post-Newtonian templates for compact binary inspiral signals in gravitational-wave detectors,” *Phys. Rev. D* **80**, 084043 (2009), arXiv:0907.0700 [gr-qc].
- [81] David Radice and Liang Dai, “Multimessenger Parameter Estimation of GW170817,” *ArXiv e-prints*, arXiv:1810.12917 (2018), arXiv:1810.12917 [astro-ph.HE].
- [82] B. Margalit and B. D. Metzger, “Constraining the Maximum Mass of Neutron Stars from Multi-messenger Observations of GW170817,” *Astrophys. J. Lett.* **850**, L19 (2017), arXiv:1710.05938 [astro-ph.HE].
- [83] L. Rezzolla, E. R. Most, and L. R. Weih, “Using Gravitational-wave Observations and Quasi-universal Relations to Constrain the Maximum Mass of Neutron Stars,” *Astrophys. J. Lett.* **852**, L25 (2018), arXiv:1711.00314 [astro-ph.HE].
- [84] M. Shibata *et al.*, “Modeling GW170817 based on numerical relativity and its implications,” *Phys. Rev. D* **96**, 123012 (2017), arXiv:1710.07579 [astro-ph.HE].
- [85] M. Ruiz, S. L. Shapiro, and A. Tsokaros, “GW170817, general relativistic magnetohydrodynamic simulations, and the neutron star maximum mass,” *Phys. Rev. D* **97**, 021501 (2018), arXiv:1711.00473 [astro-ph.HE].
- [86] Grigorios Panotopoulos and Il  dio Lopes, “Dark matter effect on realistic equation of state in neutron stars,” *Phys. Rev. D* **96**, 083004 (2017), arXiv:1709.06312 [hep-ph].
- [87] Ann Nelson, Sanjay Reddy, and Dake Zhou, “Dark halos around neutron stars and gravitational waves,” *ArXiv e-prints*, arXiv:1803.03266 (2018), arXiv:1803.03266 [hep-ph].
- [88] John Ellis, Gert H  tsi, Kristjan Kannike, Luca Marzola, Martti Raidal, and Ville Vaskonen, “Dark matter effects on neutron star properties,” *Phys. Rev. D* **97**, 123007 (2018), arXiv:1804.01418 [astro-ph.CO].

Appendix A: Complete results for simulated GW signals

In this appendix, we present complete results for the injection study described in Sec. V. In Figs. 8-12, we show the GW likelihood, non-parametric priors, and posteriors for each of the injections listed in Table I, finding the expected behavior in our posteriors. We omit process plots for the model-informed prior for all injections because they closely resemble Fig. 5. Injected signals with stiffer EOSs (**ms1**) shift the posterior to stiffer pressures relative to the prior, whereas softer EOSs (**sly**) produce lower pressures *a posteriori*. EOSs that fall near the middle of the prior (**h4**) mainly produce tightened confidence regions. We also see that the central densities for M_1 are shifted to smaller values for stiff EOSs and larger values for soft EOSs. This is because stiffer (softer) EOSs have larger (smaller) radii for the same mass, and $\rho_c \sim M/R^3$ scales inversely with the NS size. All these trends are readily apparent with the model-agnostic prior, whereas the model-informed prior often produces posteriors nearly identical to the prior. However, as Table I shows, the data typically do not strongly favor one prior over the other.

Often, **ms1** is “too stiff” for our model-informed prior but only disfavored by our model-agnostic prior. Although we attempt to construct the model-agnostic prior to support an extremely broad range of possible EOSs, there are simply many more EOSs that support relatively compact stars compared to the few that are at least as stiff as **ms1**. This is because causality places a strict upper limit on how stiff an EOS can be, evidenced by the asymmetric one-dimensional marginal distributions in, e.g., Fig. 5, and our model-informed prior assigns relatively small weight to such EOSs. The majority of tabulated EOSs used are not that stiff. Similarly, **sly** is almost “too soft” for our model-informed prior, although it still appears to be better fit than **ms1** in most scenarios. **h4** is “just right,” in that the model-informed prior captures its behavior well and is sometimes preferred over the model-agnostic one by as much as a factor of $\simeq 4$.

Generally, the model-agnostic prior’s broader support means that the *a posteriori* credible regions can stretch to include the injected values. However, we find that this does not provide a better fit to the data than the model-informed prior, with the exception of **ms1** for injection I (first row of Table I, Figs. 8 and 11). There are a few other injections for which the injected value is near the edge of the credible regions, but these are typically near the edge of the support for both the model-informed and model-agnostic priors, resulting in no strong preference for one prior over the other. We stress that this occurs because of a combination of factors. Primarily, noise fluctuations in the data shift the likelihood relative to the injected value, meaning that the maximum likelihood parameters are not the injected ones. Additionally, our priors are non-trivial in $(M_{1,2}, \Lambda_{1,2})$ and will prefer certain values over others. One should only expect credible regions to correspond to correct coverage (e.g. 50% of signals are recovered within the 50% credible regions) if the signals are actually drawn from the prior. Our injections are performed using a subset of the EOSs used to construct our priors, but they are not actually drawn from our priors. What’s more, we chose relatively extreme EOSs that live near the edge of our fiducial set, and, compounding the impact of noise fluctuations, it should consequently not be terribly surprising that some of the injected values are recovered near the edge of our 90% credible regions. Again, EOSs that are rare *a priori* are often rare *a posteriori*.




FIG. 8. Λ_1 – Λ_2 joint prior (*cyan*), likelihood (*green*), and posterior (*magenta*) distributions using the model-agnostic prior. Injected values are shown as gray cross-hairs and the colored contours correspond to 50% and 90% confidence regions. Rows correspond to events I–III from Table I and columns are labeled according to the injected EOS.

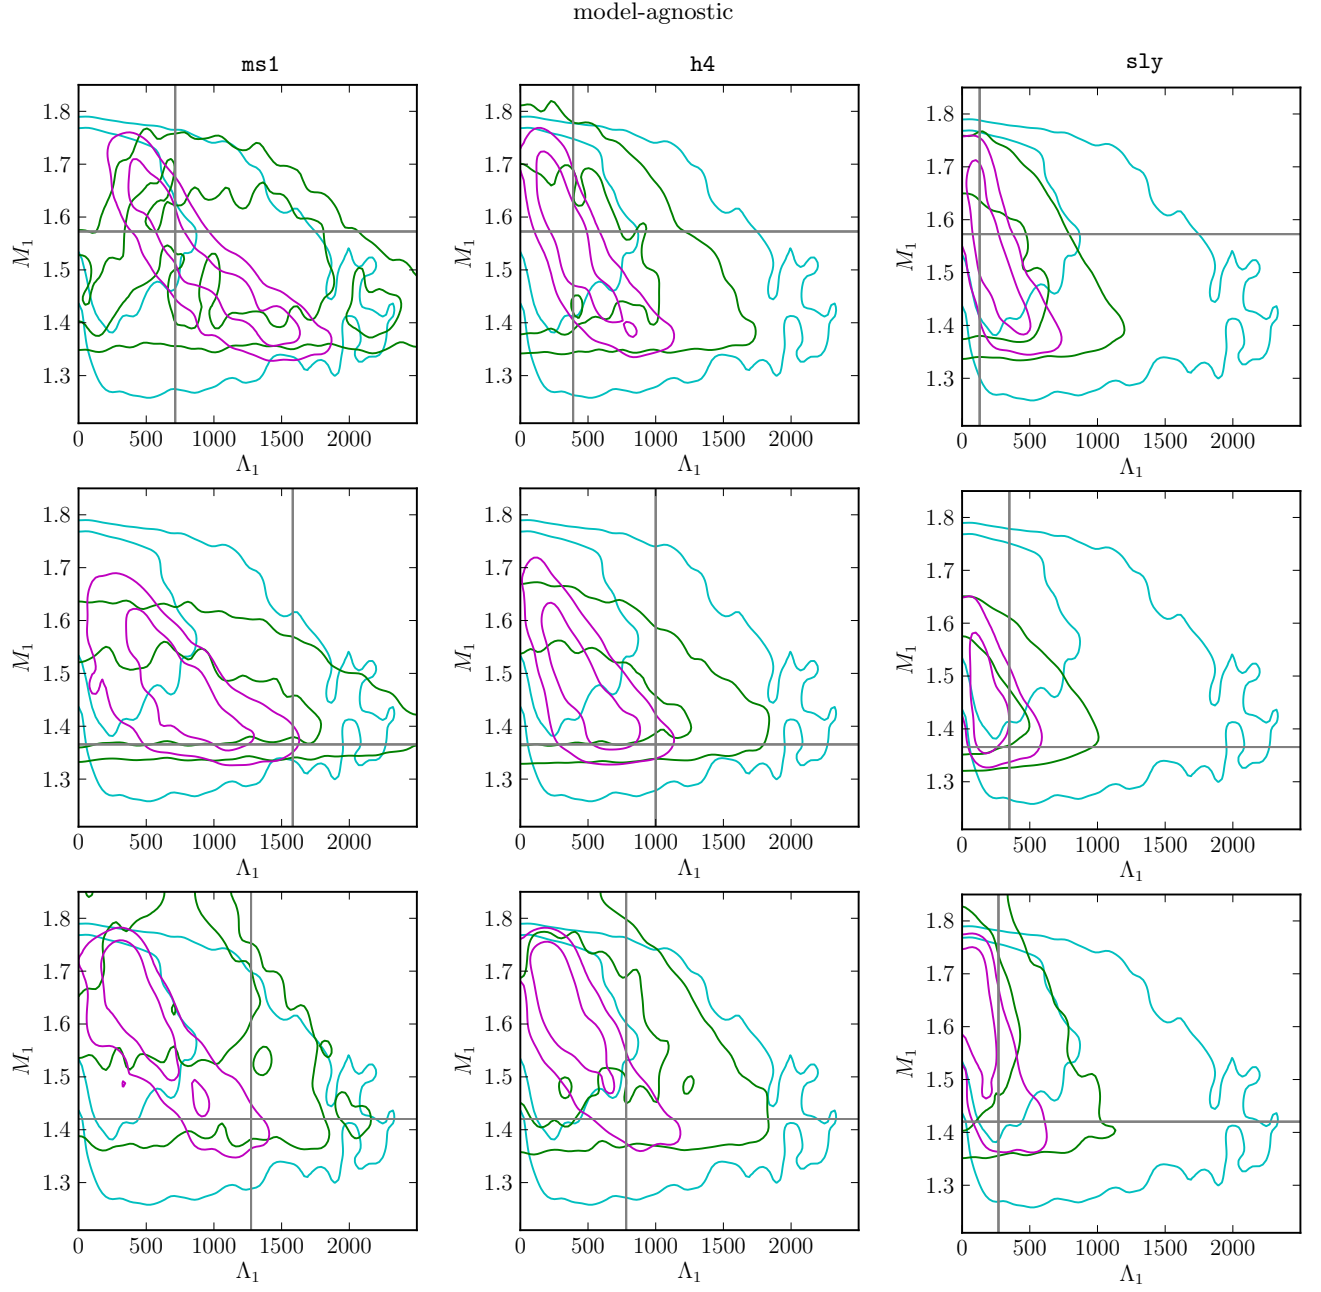


FIG. 9. Λ_1 – M_1 joint prior (*cyan*), likelihood (*green*), and posterior (*magenta*) distributions using the model-agnostic prior. Injected values are shown as gray cross-hairs and the colored contours correspond to 50% and 90% confidence regions. Rows correspond to events I–III from Table I and columns are labeled according to the injected EOS.

model-agnostic

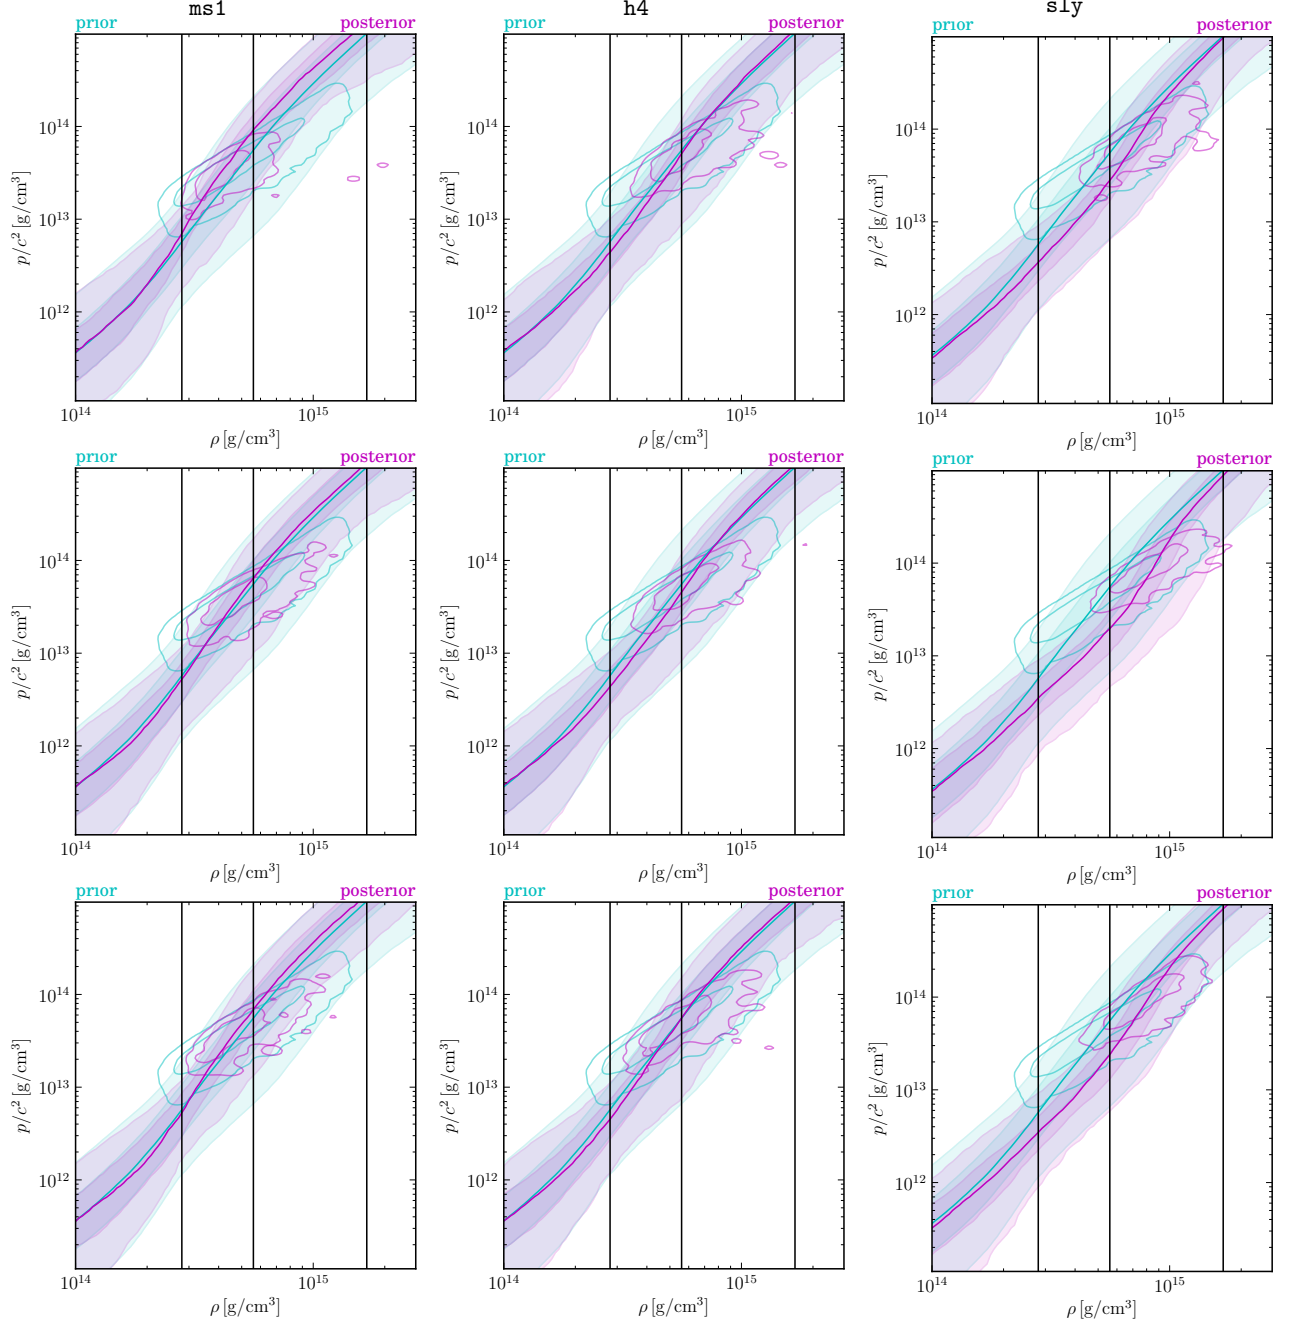


FIG. 10. Prior (cyan) and posterior (magenta) processes for injections for the model-agnostic prior. Solid lines show the medians, shaded regions represent the 50% and 90% credible regions for $p(\rho)$, and the contours represent the 50% and 90% credible regions for the central densities and pressures of M_1 . Rows correspond to events I-III from Table I and columns are labeled according to injected EOS.

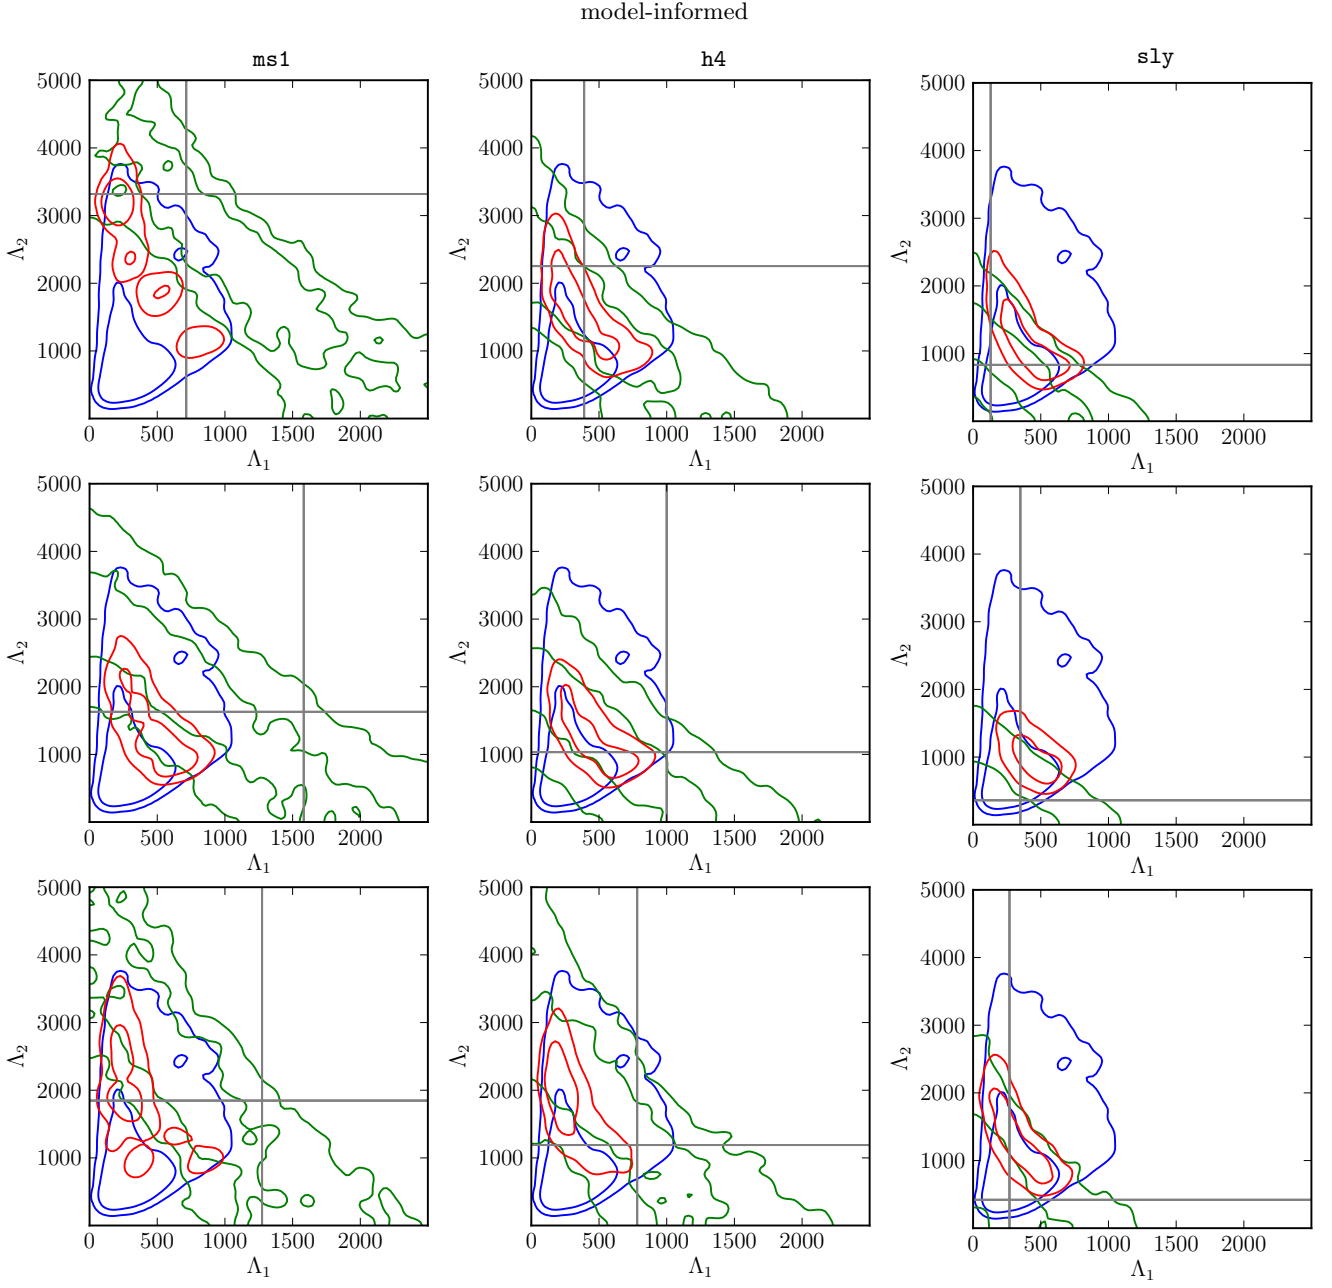


FIG. 11. Λ_1 – Λ_2 joint prior (blue), likelihood (green), and posterior (red) distributions using the model-informed prior. Injected values are shown as gray cross-hairs and the colored contours correspond to 50% and 90% confidence regions. Rows correspond to events I–III from Table I and columns are labeled according to the injected EOS. We note that **ms1**’s model-informed posterior for injection I is relatively under-sampled, with $N_{\text{eff}} \sim \text{afew} \times 10^2$ even though $N > 10^6$. What’s more, this is the only injection that strongly favors one prior over the other, with $B_T^A = 22.6 \pm 0.7$, as evidenced by the dearth of posterior support for the model-informed prior. Our evidence estimation, therefore, remains relatively precise even though errors in the posterior’s shape are dominated by the relatively small number of samples with large weights.




FIG. 12. Λ_1 - M_1 joint prior (*blue*), likelihood (*green*), and posterior (*red*) distributions using the model-informed prior. Injected values are shown as gray cross-hairs and the colored contours correspond to 50% and 90% confidence regions. Rows correspond to events I-III from Table I and columns are labeled according to the injected EOS.