

Nontrivial features in the speed of sound inside neutron stars

D. Mroczek,¹ M. C. Miller,² J. Noronha-Hostler¹ and N. Yunes¹

¹*Illinois Center for Advanced Studies of the Universe, Department of Physics, University of Illinois at Urbana-Champaign, 1110 West Green Street, Urbana, Illinois 61801-3080, USA*

²*Department of Astronomy and Joint Space-Science Institute, University of Maryland, College Park, Maryland 20742-2421 USA*



(Received 5 November 2023; accepted 22 October 2024; published 3 December 2024)

Measurements of neutron star masses, radii, and tidal deformability have direct connections to nuclear physics via the equation of state (EoS), which for the cold, catalyzed matter in neutron star cores is commonly represented as the pressure as a function of energy density. Microscopic models with exotic degrees of freedom display nontrivial structure in the speed of sound (c_s) in the form of first-order phase transitions and bumps, oscillations, and plateaus in the case of crossovers and higher-order phase transitions. We present a procedure based on Gaussian processes to generate an ensemble of EoS that include nontrivial features. Using a Bayesian analysis incorporating measurements from x-ray sources, gravitational wave observations, and perturbative QCD results, we show that these features are compatible with current constraints. We investigate the possibility of a global maximum in c_s that occurs within the densities realized in neutron stars—implying a softening of the EoS and possibly an exotic phase in the core of massive stars—and find that such a global maximum is consistent with, but not required by, current constraints.

DOI: 10.1103/PhysRevD.110.123009

I. INTRODUCTION

One of the main goals of modern nuclear physics is to determine the phase structure of QCD. The cold, catalyzed nuclear matter in neutron stars probes the zero-temperature, isospin asymmetric regime of QCD at baryon number densities (n_B) ranging from subnuclear to several times nuclear saturation density ($n_{\text{sat}} \equiv 0.16 \text{ fm}^{-3}$) in the core [1]. In this regime, first-principle QCD calculations are not yet feasible because of the fermion sign problem [2,3], and effective models and parametrizations of the equation of state (EoS) are used instead.

A variety of models have been developed to compare against astronomical observations, all of which have different regimes of validity, advantages, and disadvantages. One such model arises from chiral effective field theory (χ EFT), which breaks down at densities around $2n_{\text{sat}}$ [4,5]. In this effective theory, one prescribes a general Lagrangian that respects the symmetries of low-energy QCD (with nucleons and pions as degrees of freedom). This Lagrangian is then expanded order by order in two- and multinucleon interactions. The low-density crust ($n_B \lesssim 0.5n_{\text{sat}}$) and the high-density inner core ($n_B \gtrsim 1.1 - 1.5n_{\text{sat}}$) [6,7] of neutron stars, however, require additional modeling and assumptions beyond χ EFT about the underlying degrees of freedom and relevant interactions.

Another class of models relies on mean-field approximations of an effective Lagrangian with nucleon, electron, and muon degrees of freedom. These models lead to a

squared speed of sound $c_s^2 = dp/d\epsilon$ (where p is the pressure and ϵ is the energy density) that increases monotonically with the density (see, e.g., Refs. [8–10]). Such behavior leads to acausal sound speeds in nonrelativistic models at densities only a few times that of nuclear saturation and therefore cannot be the correct description of nuclear matter at those densities. Relativistic hadronic frameworks also break down at high densities ($n_B \gtrsim 6n_{\text{sat}}$), when nucleons start to overlap [11].

Yet another set of results is available from perturbative QCD (pQCD) calculations, in which the QCD field equations are solved perturbatively in a small-coupling expansion. These calculations have found that at very high densities ($n_B \gtrsim 40n_{\text{sat}}$) $c_s^2 \rightarrow 1/3$ (in units where the speed of light $c = 1$) from below and high-density quark matter is approximately mass-scale-invariant, or “conformal” [12–14]. Astronomical observations, however, strongly suggest that $c_s^2 > 1/3$ in the core of neutron stars [15], at densities in the range of $2 \lesssim n_B/n_{\text{sat}} \lesssim 6$. This result indicates that c_s^2 must display nonmonotonic behavior with increasing density,¹ which has motivated searches for evidence that deconfinement into approximately

¹While exact conformal symmetry implies $c_s^2 = 1/3$, and, thus, that other EoS-related quantities must take on specific values, the reverse is not a sufficient condition to establish conformal symmetry. Indeed, it is possible for c_s^2 to pass through $1/3$ a number of times before eventually approaching it from below at high densities.

conformal quark matter occurs within densities realized in neutron stars [16–19].

The onset of conformal quark matter is, however, not the only question relevant to constraining the cold nuclear EoS. Models that include heavy resonances and exotic hadronic phases and/or strange and quark degrees of freedom predict, respectively, higher-order phase transitions and crossovers and/or first-order phase transitions [1,20–40]. An N th-order phase transition occurs when the N th susceptibility of the pressure (i.e., the N th partial derivative of the pressure with respect to the chemical potential) presents nonanalytic behavior (such as a discontinuous jump or a divergence). A crossover occurs when there is no phase separation, and the change in degrees of freedom happens gradually over some range in density (i.e., all derivatives of the pressure are continuous). These different types of phase transitions and crossovers do not necessarily predict an approach of c_s^2 to 1/3 within neutron star densities (though it may happen at much higher n_B , well beyond the densities at which the star would collapse to a black hole).

Different physical processes (i.e., phase transitions of different order or crossovers) lead to unique and nontrivial structure in c_s^2 as a function of n_B (see Ref. [41] for details and extensive examples from microphysical models). Generally, a first-order phase transition is associated with the onset of new degrees of freedom. In neutron stars, a first-order transition could separate a hadronic phase from a quark phase, for example. When a first-order phase transition takes place, the presence of latent heat leads to a range in the energy density ϵ where the pressure p is constant, which appears as a *plateau* in $c_s^2(\epsilon)$ over which $c_s^2 = 0$ for a system in equilibrium. On the other hand, a crossover (or phase transitions of higher order) is typically associated with the emergence of a new state, new degrees of freedom/particles, or new interactions, that occur gradually across a range of n_B . These new particles or interactions lead to a *bump* in c_s^2 , which may be wide (like a *positive plateau*) or narrow (like a *positive spike*), depending on whether the crossover occurs over a wide or narrow region in baryon density (see examples of quarkyonic matter [24,42–44] or percolation approaches [45]). Second-order phase transitions are associated with critical points or, at vanishing temperatures, quantum critical points. In this case, c_s^2 displays a *negative spike* approaching zero (for an example at finite temperatures, see Fig. 2 from Ref. [46]). Higher-order phase transitions are also possible and may occur due to exotic baryon states or new types of interactions that could lead to a kink in c_s^2 [47]. An EoS can display one or a combination of such features depending on the assumptions made about the relevant degrees of freedom and interactions.

Recently, astronomical observations across the electromagnetic and gravitational-wave spectra have placed constraints on the macroscopic properties of neutron stars, such as the mass (M), radius (R), and tidal deformability

(Λ). These measurements have also made it possible to indirectly infer the allowed EoS via model-to-data Bayesian comparisons, since the EoS determines M , R , and Λ as a function of central number density n_B^{\max} . Analyses typically include binary tidal deformability ($\tilde{\Lambda}$) posteriors from the LIGO gravitational wave observations of events GW170817 [48–50] and GW190425 [51], the existence of heavy pulsars [52–54] and NASA’s Neutron Star Interior Composition Explorer (NICER) joint $M - R$ posteriors from PSR J0030-0451 [55] and PSR J0740 + 6620 [56] (see, respectively, Refs. [57,58] for independent analyses of these two pulsars from a separate group within the NICER Collaboration).

Other constraints are also available from the measured properties of nuclei at n_{sat} . These properties include the symmetry energy ($S = E_{\text{SNM}} - E_{\text{PNM}}$), defined as the difference in the binding energy per nucleon between symmetric nuclear matter (SNM) and pure neutron matter (PNM) as a function of density, and the slope parameter (L), which determines how the symmetry energy changes with density [59–65].

From the theory perspective, it recently became possible to consistently extrapolate pQCD results to densities as low as $\sim 2.5n_{\text{sat}}$ [66,67]. These constraints are based on the mechanical stability and causality of the EoS ($0 \leq c_s^2 \leq 1$) and the consistency of the underlying thermodynamic potential that connects the low-density regime of the EoS to the high-density regime ($\gtrsim 40n_{\text{sat}}$) constrained by pQCD. These constraints offer information at each n_B about the region in $p - \epsilon$ that can be connected to the high-density perturbative results via a stable and causal EoS and a consistent thermodynamic potential. For a given EoS, it is possible to check its compatibility with stability, causality, and consistency constraints at any density between $\sim 2.5 - 40n_{\text{sat}}$. We will refer to these constraints collectively as the pQCD constraints from here on.

The lack of first-principle approaches for the β -equilibrated, zero-temperature nuclear EoS between $\sim 1.1n_{\text{sat}}$ up to n_B^{\max} realized in neutron stars means that astronomical observations are the only direct probe of the EoS in this regime. Thus, model-to-data Bayesian comparisons of generic functional forms of the EoS are the state of the art for obtaining posterior distributions for the EoS. Nonetheless, microscopic models are vital in providing guidance for the behavior of functional forms of the EoS, especially so that specific features associated with the onset of new degrees of freedom and interactions can be correctly identified.

The posterior distribution that is extracted from a Bayesian analysis is sensitive to how data and theoretical input are incorporated [68] as well as prior-imposed assumptions about the EoS (e.g., correlations across density scales) [69–72]. Parametric descriptions, such as spectral expansions [73,74] or piecewise polytropes, provide a framework to represent the EoS without relying on

microphysics models. Spectral representations of the EoS assume the adiabatic index $\Gamma(p)$ as a function of pressure can be expanded in terms of a set of spectral basis functions and coefficients, which uniquely determine the EoS [73–75]. Piecewise polytropes divide the EoS into a number of segments and represent the pressure for each segment as a polytrope, $p = \kappa \rho^\Gamma$, where κ and Γ are the fixed polytropic constant and the adiabatic index, respectively [76]. Parametric representations have been widely used, since they do not rely on as many assumptions as physics-based models (see, e.g., Refs. [16–18,50,55,56,76–82], though this list is far from comprehensive). However, the question of whether these parametrizations are flexible enough to capture all relevant physics has recently been raised in the literature [69,72,83]. Specifically, it has been shown that parametrized EoS can introduce undesired correlations across density scales [69,72] and are unable to capture behavior consistent with state-of-the-art nuclear physics models with exotic degrees of freedom [83].

Physics-agnostic frameworks based on Gaussian processes (GPs) offer more flexibility in the modeling of the EoS at the cost of increased functional complexity [15,84–87]. Generally, a GP models the speed of sound as a continuous function over a specified domain. The properties of the probability density for the speed of sound at each point of the domain are determined by a mean vector and covariance matrix. The covariance matrix is calculated using a specific kernel function that requires hyperparameters that can be fixed or sampled from a hyperprior. The hyperprior may be model agnostic or conditioned to more closely reproduce a set nuclear physics models [87]. In principle, a GP can be tailored to resemble any continuous function across some domain. So far, GPs have been implemented with a fixed set of hyperparameters for an individual EoS [15,16,56,58,67,84–87], though priors may contain samples drawn from mixture of multiple stationary kernels, probing a wide range of potential correlation properties [15,87]. In contrast with the assumption of uniform correlations across density scales, many state-of-the-art nuclear physics models with exotic degrees of freedom display multiscale correlations across various densities [83]. Furthermore, the features that emerge in the speed of sound as a result are known to be important for understanding heavy, ultraheavy (neutron stars with masses above $2.5M_\odot$), and twin stars [41,83,88].

With that motivation, we introduce modified GPs (mGPs) as a framework for modeling EoS with nontrivial features that possess long-, medium-, and short-range correlations across densities. First, we produce a family of EoS from a benchmark model of GP EoS that contain only long- and fixed-range correlations in c_s^2 . We then generate a family of EoS from mGPs, which introduce multiscale correlations in the form of nontrivial features in c_s^2 . With these two families of EoS, we carry out Bayesian parameter estimation analysis against observational and experimental data and

input from pQCD. The results of this analysis allow us to compare the marginalized posteriors of the mass and radius curve and the speed of sound and number density curve when we use the benchmark GP model and the modified GP to represent the EoS. We find that neither EoS family is favored over the other by current data. We do find, nonetheless, that the marginalized posterior for the speed of sound at densities $\sim 1.5–2n_{\text{sat}}$ is not identical for EoS from GPs compared to mGPs, although the data are not informative enough yet to discriminate between these posteriors.

The remainder of this manuscript presents the details of the analysis summarized above. In Sec. II, we discuss GPs as a model-agnostic framework for generating the EoS and how we introduce multiscale correlations to the EoS with mGPs. Section III outlines how we generate EoS priors from the GP and mGP frameworks. Statistical methods are discussed in detail in Sec. IV, followed by results in Sec. V and conclusions and discussion in Sec. VI. Throughout this manuscript, we use $c = 1$ and the Einstein summation convention when necessary. Thermodynamic quantities are in centimeter-gram-second (cgs) units (unless otherwise stated), with the exception of c_s^2 , which we always normalize by c^2 .

II. GENERATING THE EQUATION OF STATE

Both the benchmark EoS model that is a standard GP and the modified GP EoS are built from model-agnostic GPs, which approximate functional forms of c_s^2 as a function of the pressure over a fixed domain (for a more comprehensive overview of Gaussian processes, we recommend Refs. [84,89].) We now discuss the details and motivation for the construction of both models.

A GP provides the joint probability density for a continuous function $f(x)$ over a domain of interest, which we represent here by a sample of discrete values labeled x_i . This probability density is assumed to be a multivariate Gaussian distribution (no summation over i implied),

$$f(x_i) = \mathcal{N}[\mu_i(x_i), \Sigma_{ij}(x_i)], \quad (1)$$

where $\mathcal{N}(\cdot, \cdot)$ is the normal distribution function at x_i , with a mean μ_i that varies with x_i and a covariance matrix Σ_{ij} , which gives the correlation between the values of f at x_i and x_j , where i can equal j .

In the context of extracting the properties of neutron stars from data using an ensemble of synthetic EoS, GPs have been used to approximate the EoS from samples of functional forms of c_s^2 [15,84,86,87,90–92]. Because the range of Gaussian distributions is infinite, while c_s^2 is bounded by stability and causality ($0 \leq c_s^2 \leq 1$), it is common to use the GP to approximate an auxiliary variable that compactifies the range of a Gaussian distribution (infinite in both the positive and negative directions) to the range of c_s^2 . Let us call this variable ϕ [73,84] and define it via

$$\phi \equiv \ln(de/dp - 1) = \ln(1/c_s^2 - 1), \quad (2)$$

where p is the pressure and ϵ is the energy density. This auxiliary variable has the desired range for a GP, but when mapped to c_s^2 using the definition above, $\phi \rightarrow +\infty$ corresponds to $c_s^2 \rightarrow 0$, while $\phi \rightarrow -\infty$ corresponds to $c_s^2 \rightarrow 1$. It is common in the literature to model ϕ as a function of $\log_{10} p$ in cgs units [15,56,84,86,87,91,92]. More explicitly, Eq. (1) becomes

$$\phi(\log_{10} p_i) = \mathcal{N}[\mu_i(\log_{10} p_i), \Sigma_{ij}]. \quad (3)$$

Other units and thermodynamic variables can be used instead (e.g., Ref. [90] uses baryon density in units of n_{sat}), but we will use the parametrization presented above.

A more computationally practical implementation of a GP is to decompose it into a mean and a scatter via

$$\phi(\log_{10} p_i) = \mu_i(\log_{10} p_i) + L_{ij} u_j, \quad (4)$$

where L_{ij} is the Cholesky decomposition of the covariance matrix plus a white-noise kernel contribution, i.e., $L_{ik} L_{kj}^T := \Sigma_{ij} + \sigma_{\text{wn}}^2 \delta(x_i - x_j)$, with $u_j = \mathcal{N}(0, 1)$ and σ_{wn} a constant white-noise variance. The white-noise kernel (the second term on the right-hand side of the Cholesky decomposition) is added for numerical stability, since the determinant of Σ_{ij} can be nearly singular. The effect of the white-noise kernel is to slightly smear the GP by adding noise to the diagonal elements of Σ_{ij} . A small σ_{wn} is sufficient to dramatically increase the stability of the calculation without changing the overall properties of the final sample. We use $\sigma_{\text{wn}}^2 = 0.0003$, but any other value of the same order of magnitude would produce similar results.

Given the joint probability density for ϕ , we can construct a GP realization, or “sample,” by selecting a range of pressures and then drawing the associated ϕ values using Eq. (4) [64]. We then invert Eq. (2) to find c_s^2 as a function of p . The definition of the speed of sound can then be used to specify a differential equation for the EoS, $c_s^2(p) = dp/d\epsilon$, which can be solved in first quadrature as

$$\epsilon = \int \frac{dp}{c_s^2(p)}, \quad (5)$$

and then inverted to find $p(\epsilon)$. We then obtain the baryon density using the first law of thermodynamics, which, at zero temperature and assuming charge neutrality, can be written as

$$\frac{dn_B}{d\epsilon} = \frac{n_B}{\epsilon + p(\epsilon)}. \quad (6)$$

Once these equations are solved, we have the set $\{\phi(p), c_s^2(n_B), p(\epsilon)\}$, which defines an EoS sample from

a GP. In practice, we build our EoS numerically by sampling on a finite set of pressures and baryon densities with a sufficiently fine discretization. As pointed out in Ref. [41], a simple check that the EoS is being recovered correctly is to calculate c_s^2 from p and the reconstructed ϵ and check that it matches the c_s^2 from the GP. The EoS samples generated through the GPs will only be used above $0.5n_{\text{sat}}$ [80], which we denote as the core-crust transition. Below this density, we model the crust through the QHC19 EoS [39,93].

A. Benchmark Gaussian processes

Now that we have explained the idea behind constructing an EoS sample from a GP, we need to specify the input for the joint probability density function in Eq. (3). The two main ingredients that define a GP are the *means* $\{\mu_i\}$, which will determine the average trend for the function that is being sampled, and the *covariance matrix* Σ_{ij} , which specifies the joint variability between two points x_i and x_j .

Let us first discuss how we model the means. Our goal at this stage is to create a benchmark model $c_s^2(p)$ *without* any sharp, nontrivial features. To do so, we adopt the approach taken in Miller *et al.* [56], which looked at a collection of 12 cold neutron star EoS from the CompOSE database [94,95] on the $\log_{10} p - \phi$ plane, and find that these EoS follow a linear trend over the domain $32.7 \leq \log_{10} p_i(\text{erg cm}^{-3}) \leq 37$. This trend was empirically approximated as

$$\mu_i(\log_{10} p_i) = 5.5 - m(\log_{10} p_i - 32.7), \quad (7)$$

where m is the slope of the linear regression. Reference [56] fixes $m = 2$ based on the spread of EoS from CompOSE. Other choices for the means are also possible [87]. Out of the total 12 EoS that this model is based on, seven were purely proton, neutron, and electron matter (npe) models [8,9,96–104]; one model included npe matter, heavy baryonic resonances, and a crossover transition to quarks (QHC18 [1,93,105]); and four models included npe matter plus strange baryons [8,102–104,106–110]. These models largely approach the causal limit at high densities, which biases the behavior of the EoS in that regime. Notably, models that predict a softer EoS at large densities, such as quarkyonic models [24,42–45], are missing from the collection of EoS that was used to determine Eq. (7).

To test the assumption of Miller *et al.* [56], we use the relation in Eq. (7) with $m = 2$ to create a set of EoS samples. As shown in the top left panel of Fig. 1 (solid, thin, gray lines), these EoS samples cluster around a mean (solid, thick, black line). The top right panel of this figure shows that speed of sound functional forms constructed using $m = 2$ largely approach unity with increasing density. This behavior is highlighted in the bottom panel, which shows c_s^2 as a function of pressure. Note that different EoS have different ranges in baryon density for the same range of

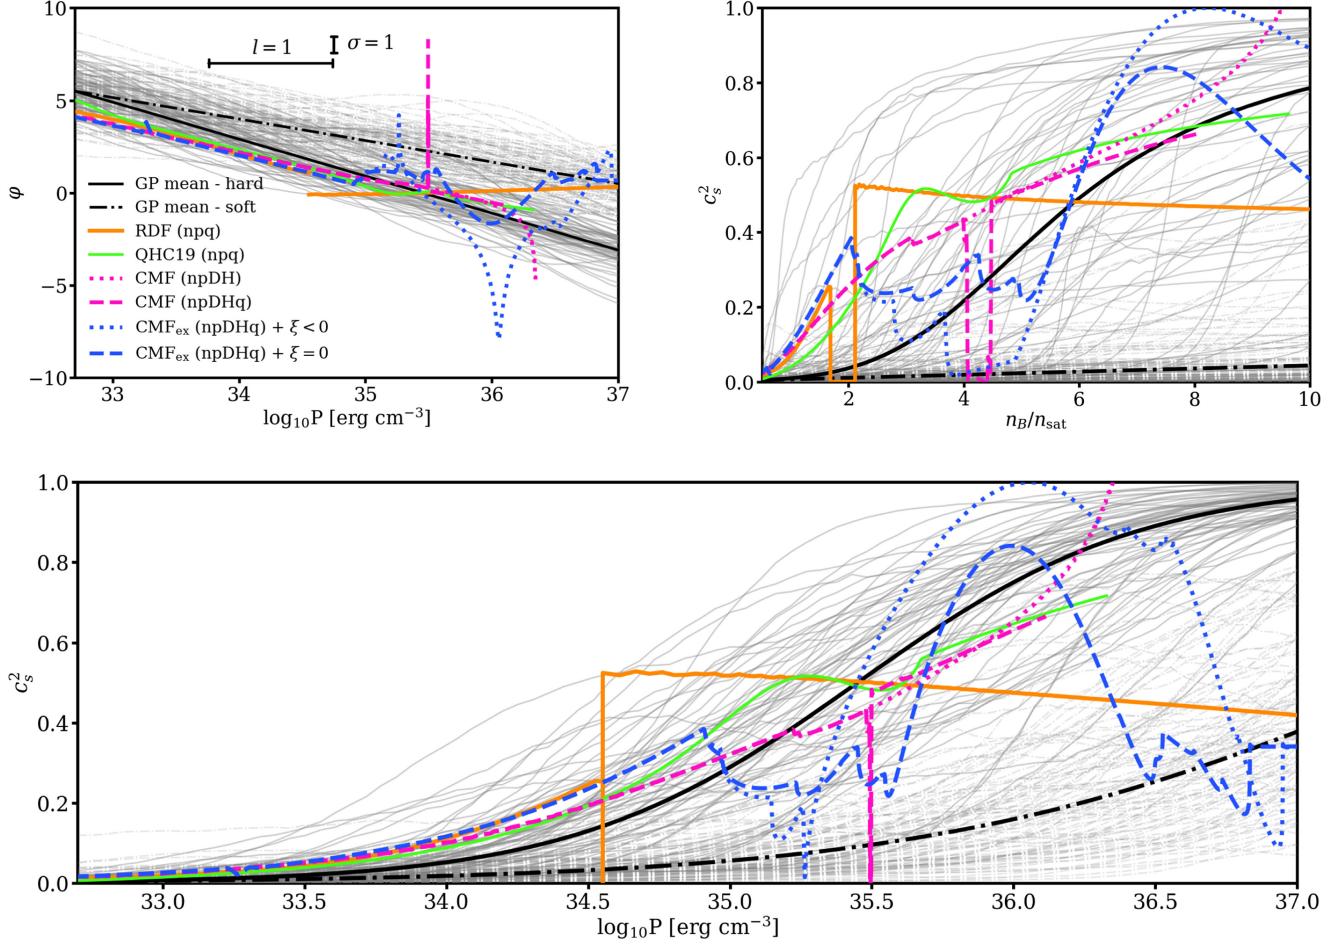


FIG. 1. The auxiliary variable $\phi \equiv \ln((de/dP) - 1)$ as a function of $\log_{10} P$ in cgs units (left) and the speed of sound as a function of baryon number density in units of n_{sat} (right) for different parametrizations of the EoS from CMF, which include nucleons (np), D, H, quarks (q), and leptons [109–115], the QHC19 EoS (np + q + leptons) [39], and the relativistic density-functional (RDF) approach introduced in Ref. [116] (np + q + leptons). In light gray, we show a sample of 160 total functional forms for c_s^2 from the benchmark GPs using the ansatz in Eq. (7) and a squared-exponential kernel with $l = \sigma = 1$. The functional forms from the hard and soft GP are shown in dot-dashed and solid lines, respectively. The benchmark GPs capture a wide range of behavior, but the *a priori* requirement that functional forms display only long-range correlations across densities exponentially suppresses sharp and nontrivial features in c_s^2 , which are observed in state-of-the-art nuclear physics simulations.

pressure, so the range of pressures in the top left panel does not correspond to the range of baryon densities in the top right panel. We now contrast this set of EoS samples with a new set, constructed from GPs with a softer mean. More specifically, we set $m = 1.6$ in Eq. (7), resulting in the functional forms shown in the top left panel of Fig. 1 (dot-dashed, thin, gray lines). As expected by construction, the mean of these samples has a softer slope (dot-dashed, thick, black line). The effect of this softer mean is to reduce the speed of sound to values largely below ≈ 0.4 in the neutron star range of baryon densities, as shown in the top right panel of Fig. 1. For baryon densities larger than what we expect in neutron stars, the distribution of speeds of sound has a mean of $1/3$ (i.e., the conformal limit), and a scatter that leads to c_s^2 's as large as 0.8 and as small as 0.1, as shown in the bottom panel. From here on, we refer to the set of EoS

samples resulting from GPs with $m = 2$ and $m = 1.6$ as “hard GP” and “soft GP,” respectively. Figure 1 also presents specific realizations of nuclear physics simulations of the EoS, but we defer a discussion of those to Sec. II B.

Why consider a soft GP when astronomical observations seem to indicate that the conformal limit is broken at $n_B \approx 2n_{\text{sat}}$ [7,15,56,117]? Our motivation is to show the effect of softer means in the speed of sound, while at the same time generating a new benchmark model that can be modified through sharp features in a narrow baryon density range to make them consistent with astronomical observations. We will discuss such modifications in the next subsection.

Before proceeding, let us discuss two other important modifications from the approach in Miller *et al.* [56]. The highest value sampled in $\log_{10} p$ is 37, instead of 36 in Ref. [56], and we use a significantly finer grid—Miller

et al. samples 50 points, whereas we sample 100. These modifications are necessary because our procedure allows for softer EoS, which result in higher neutron star central pressures. Expanding the sampled domain ensures that the entire stable branch is captured, rather than cutting it off at an arbitrary, smaller value. Also, since in the next subsection we will introduce sharp features in c_s^2 that lead to rapid changes in the EoS, a finer grid is needed to keep numerical errors under control when recovering the EoS samples from a GP.

Let us now discuss the second ingredient that defines a GP: the covariance matrix. We assume that Σ_{ij} is a matrix whose elements are determined through a kernel function of the pair $\{x_i, x_j\}$, where x_i is the point at which the normal distribution is being sampled (i.e., in our case, x is the \log_{10} of the pressure in cgs units) and x_j is any other point, i.e., $\Sigma_{ij} = K(x_i, x_j)$. We further assume a squared-exponential kernel,

$$K_{\text{se}}(x_i, x_j) = \sigma^2 \exp \left[-\frac{(x_i - x_j)^2}{2\ell^2} \right], \quad (8)$$

which depends only on the distance between x_i and x_j and on two hyperparameters, ℓ and σ . Specifically, ℓ determines the correlation length scale (e.g., when $\ell \rightarrow 0$ all points are independent of each other), and σ represents the strength of the overall correlation (e.g., when $\sigma \rightarrow 0$ all points go to the mean). The benchmark models should be smooth, meaning that ℓ should be compatible with longer correlations across domain points. In addition, because of the exponential map between ϕ and c_s^2 , a σ that is too large would lead to $c_s^2 \approx 0$ or $c_s^2 \approx 1$ more often. Nonetheless, σ should not be too small, so there is enough variability in the EoS samples from any given GP. In accordance with Miller *et al.* [56], we set $\ell = \sigma = 1$ for both the hard and soft GP benchmark models.

Let us now consider the effect of our choice of ℓ and σ on our EoS samples. Figure 1 shows that, in the hard GP case, $\sigma = 1$ still allows for enough deviation from the mean to create variability in the EoS samples, without oversampling $c_s^2 \approx 1$ or $c_s^2 \approx 0$. In the case of the soft GP, a $\sigma = 1$ leads to oversampling $c_s^2 \approx 0$, since the mean is already at very low values of c_s^2 . However, for both GPs, $\ell = 1$ heavily suppresses large deviations in c_s^2 from one value of pressure to pressures in a close neighborhood, resulting in EoS without sharp, nontrivial features. Our assumptions in the benchmark models do not force c_s^2 to increase monotonically; nevertheless, because $\ell = 1$ imposes large-scale correlations, nonmonotonic behavior is smeared out across a wide range of densities, a feature that is consistent only with a smooth (i.e., wide) crossover.

B. Modified Gaussian process

Are the benchmark models discussed in the previous subsection enough to accurately represent nuclear-physics-derived EoS? Figure 1 shows a set of EoS derived from state-of-the-art simulations of chiral mean field theory (CMF) models [109–114], a simulation of the commonly used quark-hadron crossover EoS framework (QHC19) [39], and one example from the RDF model with density-dependent vector and diquark couplings [116,118]. In particular, we include in this figure CMF models with delta resonances (D), hyperons (H), and quark (q) degrees of freedom, where the transition to quark degrees of freedom is a first-order phase transition (denoted CMF) or a crossover due to an excluded volume term (denoted CMF_{ex}) for two different parametrizations of the strange vector quark couplings [115]. The RDF example included here corresponds to the onset of a two-flavor color-superconducting quark phase at $n_B = 0.287$ 1/fm³, the central density of a $\sim 1M_\odot$ star in this model. As shown in Fig. 1, exotic degrees of freedom lead to kinks, spikes, and plateaus in c_s^2 that occur across short correlation lengths in baryon density.² None of the EoS samples drawn from either of the two benchmark GPs is able to reproduce these features.

This discrepancy between the benchmark GPs and nuclear physics simulations motivates the creation of mGP. More specifically, we wish to create a modification to the benchmark GPs that lead to EoS samples that contain the short-length correlation structures in the speed of sound that are present in realistic nuclear-physics simulations, while maintaining long-length correlation scales driven by an overarching mean behavior. A mGP sample is built from a benchmark GP that serves as a *baseline* but that is *modified* through the addition of a specific feature in a range of pressures. We do not introduce modifications below n_{sat} because a variety of experimental constraints (see Ref. [60,119] for a recent review) and χ EFT calculations [5] require no such structure at these low densities.

Two main reasons drive our choice to introduce modifications to a baseline GP. The first is direct control over the functional form of c_s^2 at a low computational cost. Each modification that is introduced can be directly related to a thermodynamic process, and we have the ability to track where and how modifications appear without any post-processing. The second is *a priori* multiscale correlations in density. We note that GPs with a fixed, but sufficiently small correlation-length can converge to an EoS that displays long, medium, and short correlations in density *a posteriori*, in which case convergence (i.e., the posterior credible regions are small and centered around the true EoS at all density/mass scales) may require a large number of

²See also Fig. 2 in Ref. [83] for more examples of nuclear physics simulations of EoS with exotic degrees of freedom and how nontrivial features appear in the speed of sound.

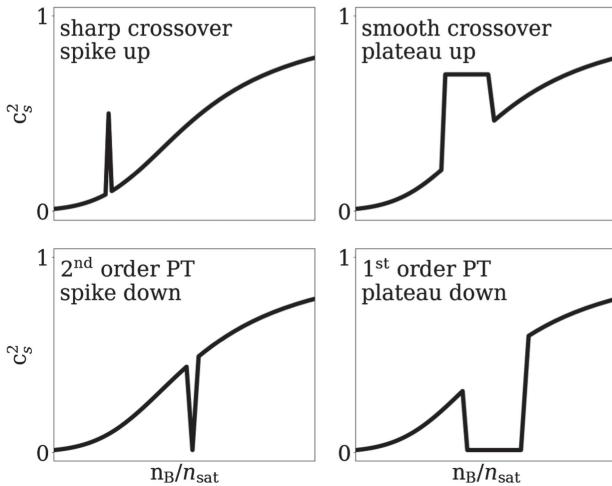


FIG. 2. Drawing of c_s^2 as a function of n_B to illustrate different physical features to incorporate in mGPs. A sharp crossover corresponds to a rapid change in degrees of freedom, where c_s^2 will first become stiffer—due to the onset of, e.g., repulsive or excluded volume interactions—and then quickly soften with the emergence of new degrees of freedom; this leads to a spike/sharp bump upward in c_s^2 with respect to the baseline (top left). A smooth crossover corresponds to a slower change in degrees of freedom, which we model as a plateau upward with respect to the baseline (top right). A second-order phase transition corresponds to the critical point at the end of a first-order phase transition line, leading to a very small region (approximately a *point*) where $c_s^2 = 0$, which can be modeled as a spike downward with respect to the baseline (bottom left). Last, a first-order phase transition separates two phases with distinct densities, leading to a gap in c_s^2 , which can be represented as a plateau, or region, where $c_s^2 = 0$ (bottom right).

samples.³ However, if the GPs are constrained to larger correlation lengths *a priori*, medium- and short-range correlations will be exponentially suppressed, and an even larger number of samples would be required to converge to a posterior that displays multiscale correlations [120].

In the following paragraphs, we connect the types of modifications we introduce in the mGPs to the phase transition phenomenology from numerical simulations of nuclear physics models and general thermodynamic arguments.

There are a few different thermodynamic phenomena relevant to study of phase transitions. These phenomena are illustrated in Fig. 2 as functional forms of the speed of sound as a function of baryon number density and Table I,

³It is also important to note that data may not constrain large changes in the speed of sound over a short range in density very strongly. That is, the likelihood for any individual event may not be very informative, which would required many events to get an informative joint likelihood. In that case, any tighter credible regions derived with priors constrained to display short-range correlations would be due to the prior rather than the likelihood [120].

which connects the different features illustrated in Fig. 2 to relevant physical processes and nuclear physics models. We will now discuss each of the categories of physical phenomena and modifications individually and draw connections to Fig. 2 and Table I.

In general, a phase transition can be continuous (also known as a *crossover*) or discontinuous. If continuous, all derivatives of the pressure are finite, i.e.,

$$\left(\frac{\partial^n p}{\partial \mu_B^n}\right)_{\text{crossover}} \neq \infty, \quad (9)$$

and the pressure is an analytic function of the energy density. If a discontinuity exists, i.e.,

$$\left(\frac{\partial^n p}{\partial \mu_B^n}\right)_{\text{nth-order}} \rightarrow \infty, \quad (10)$$

the phase transition is classified through its order n . An (n th)-order phase transition is one in which the (n th)-order derivative of the pressure with respect to the chemical potential at the critical baryon density $n_B = n_{\text{crit}}$.

At vanishing temperatures, crossovers lead to a non-monotonic peaklike behavior—i.e., a bump—in the speed of sound (see Ref. [24] for an example and explanation of this behavior). Mathematically, we can define the center of this bump at \tilde{n}_B , and half the range in n_B where this bump occurs will be defined as δ . Then, if the peak of the bump is centered at \tilde{n}_B (we make the assumption here that \tilde{n}_B is centered at the middle of the bump but it does not necessarily have to be the case), we find that at $n_B = \tilde{n}_B$, $(c_s^2)' \sim 0$.

Crossovers can be further classified into *smooth* (see, e.g., Ref. [24]) or *sharp* (see, e.g., Ref. [121]), depending on the abruptness of the nonmonotonic behavior. A sharp crossover is when the change in degrees of freedom happens rapidly (as expected in quarkyonic matter models [24,42–44,121], for instance), such that $\delta \ll 1n_{\text{sat}}$ and the peak behavior becomes more of a “spike” (although all derivatives of the pressure remain finite). This kind of crossover is summarized in the top row of Table I, is illustrated through the drawing in the top left panel of Fig. 2, and will be represented in the mGP framework by a spike that rises relative to the baseline. A smooth crossover is when the change in degrees of freedom happens slowly so that $\delta > 0$ is large and the peak behavior becomes more of a “plateau” such that there is no longer a single sharp point in baryon density where the derivative of c_s^2 is zero but rather a range of $n_B \sim \tilde{n}_B$ such that $(c_s^2)' \sim 0$. This kind of crossover is summarized in the second row of Table I, is illustrated through the cartoon in the top right panel of Fig. 2, and will be represented in the mGP framework by constant $c_s^2 \neq 0$ in a $\tilde{n}_B \pm \delta$ region. For both sharp and smooth bumps, we define the bump as an increase of at

TABLE I. Connection between phase transitions of different orders/crossover to corresponding physical processes in terms of the effect on the speed of sound in equilibrium and modifications in the mGP framework. Note that a first-order phase transition has a jump in baryon density across Δn_B .

Transition type	Physical process	Representation in c_s^2	Modification
Sharp crossover	Quarkyonic matter [24,42–44,121], percolation to quark matter [39,45], quark-meson coupling [35,122] heavy resonances [109–115], hyperons [47,109,123], chiral-superfluid transition [124]	For $\delta \ll 1$, if $n_B = \tilde{n}_B \pm \delta$, then $(c_s^2)' = \pm \delta^{-1}$ if $n_B = \tilde{n}_B$, then $(c_s^2)' = 0$	Spike up, $c_s^2 \neq 0$
Smooth crossover		For $\delta > 0$, if $n_B = \tilde{n}_B \pm \delta$, then $(c_s^2)' = \pm \delta^{-1}$ if $n_B \sim \tilde{n}_B$, then $(c_s^2)' \sim 0$	Plateau up, $c_s^2 = 0$
n th-order PT, $n > 2$		If $n_B = n_{\text{crit}}$, then $d^n p / d\mu_B^n \rightarrow \infty$	Spike or plateau down, $c_s^2 \neq 0$
Second-order PT	Critical point due to exotic quark phases	$c_s^2(n_{\text{crit}}) = 0$	Spike down toward $c_s^2 \approx 0$
First-order PT	Quark deconfinement [109,111], color superconductivity [116,118], color-flavor locking [26]	$c_s^2(n_B) = 0$ with $n_B \in [n_B^*, n_B^* + \Delta n_B]$	Plateau down at $c_s^2 \approx 0$

least 10% compared to the original benchmark functional sampled in the regime of $\tilde{n}_B \pm \delta$.

Let us consider an example of a sharp crossover in more detail, taking the quarkyonic model as a reference [24,121]. In this framework, the speed of sound squared is always below the conformal value of 1/3, except in a narrow range of densities where the crossover transition happens. The rapid stiffening is associated with repulsive, excluded volume interactions, followed by a softening of the speed of sound, once quark and gluon degrees of freedom appear in the system. A scenario like this one is equivalent to a soft GP baseline with a spike that rises up in a small baryon density region. Our mGP also includes more general cases where a spike up is added to a hard GP baseline, meaning that c_s^2 will not be required to stay below the conformal value neither before nor after the crossover, since those cases cannot yet be ruled out by the data.

Let us now consider a few examples of smooth crossovers in more detail. A minimal set of requirements to create a plateau or “bump” structure is discussed in Ref. [124], which found that in QCD this feature can be the result of a “chiral-superfluid” transition, such as the condensation of diquarks or dibaryons. In quarkyonic frameworks, a plateau structure can appear when repulsive excluded-volume terms are partially balanced by the onset of quark degrees of freedom, which stiffen and soften the EoS, respectively [42–44]. In the CMF model, implementing an excluded-volume term for the hadrons leads to a crossover transition to the quark phase [115]. In the quark-hadron crossover EoS (QHC) [39], or three-window modeling of the EoS [45], the crossover regime is constructed via a smooth interpolation of the hadronic and quark regimes, which must also respect thermodynamic

constraints such as causality. Lastly, the quark-meson coupling (QMC) class of models, which is based on baryons that interact via the exchange of virtual mesons between confined valence quarks, also gives rise to smooth crossover structure in c_s^2 [35,122]. Furthermore, QMC EoS soften rapidly with the onset of hyperons, leading to $c_s^2 < 1/3$ within neutron star densities even when no quarks are produced [35], a feature that is relevant for recent discussions on the onset of a conformal regime in the core of massive neutron stars [16,18,67]. Scenarios like the ones described above lead to a rounded peak structure in the speed of sound (see Fig. 1), which can be approximated as a plateau at some finite c_s^2 that rises above the baseline EoS. We make this approximation for simplicity, given that there is an infinite number of continuous functions that can be constructed in the crossover regime. Although we do not expect this approximation to affect macroscopic observables significantly, it would be valuable to quantify its impact in a future study, accounting for variables such as the width and height of the peak and density dependence. A smooth crossover can also be constructed phenomenologically (see, e.g., Refs. [125–127]).

Discontinuous phase transitions of order higher than 2 lead to speeds of sound that resemble that of crossovers, although technically the higher derivatives of the pressure are not defined and the pressure is thus a nonanalytic function of energy density. In fact, for some models, it is still an open question whether certain phase transitions are crossovers or discontinuous phase transitions of finite order [47]. For this reason, we will model discontinuous phase transitions of order higher than 2 through spikes and plateaus in the speed of sound that dip below the baseline but do not lead to vanishing c_s^2 . The inclusion of these

features leads to a variety of nonmonotonic structure in c_s^2 across some finite range in n_B .

A discontinuous phase transition of order 2 is sometimes referred to as a “critical point” (illustrated in Fig. 2, bottom left) because it is the end point of a first-order phase transition line. As presented in Table I, we are not aware of any models that predict a critical point at zero temperature for β -equilibrated nuclear matter (note that a zero temperature critical point is known as a quantum critical point). However, this possibility cannot currently be ruled out by the data, and thus we choose to model it. At a critical point, $c_s^2 \approx 0$ only at a critical baryon density n_{crit} . In the mGP framework, we model a critical point as a spike that dips to low values of c_s^2 (below 0.1) at a single value of baryon density. Ideally, we want $c_s^2(n_{\text{crit}}) = 0$, but because we reconstruct the EoS numerically, taking $c_s^2(p)$ to be exactly zero at any point along the grid would cause the expression for the energy density to become nonanalytic. Additionally, because we add a modification to a benchmark EoS that has the same correlation length before and after the transition, large deviations in the value of c_s^2 before and after the transition are suppressed in the prior. However, these values could in principle be completely uncorrelated. In future work, our model could be improved by considering different benchmark EoS before and after the transition. Nonetheless, the changes introduced by a spike going to low c_s^2 are significant on the $\epsilon - p$ plane (see Appendix).

A discontinuous phase transition of order 1, also known as a *first-order phase transition*, occurs when the transition between two different phases of matter requires a nonzero latent heat. As a result, the two phases have different baryon densities. If the transition density is n_B^* , then the speed of sound displays a gap, i.e., a region where $c_s^2 \approx 0$, between n_B^* and $n_B^* + \Delta n_B^*$. The larger the gap, the stronger the phase transition. This description assumes that the system is in equilibrium and a Maxwell construction was performed to remove any metastable region. In a dynamical system, the first-order phase transition would present as a metastable, or spinodal, region wherein one would see nonmonotonic behavior in $c_s^2(n_B)$, $n_B \in [n_B^*, n_B^* + \Delta n_B^*]$. Since neutron stars are in equilibrium, a Maxwell construction across the phase transition is a good assumption, leading to a plateau in $p(\epsilon)$ that results in a region of $c_s^2(n_B) = 0$, $n_B \in [n_B^*, n_B^* + \Delta n_B^*]$.

First-order phase transitions arise in a variety of nuclear-physics models, as presented in Table I. In the CMF framework [109,111], a first-order phase transition results from a Polyakov loop being used to describe the separation between the hadronic phase with deltas and hyperons from the quark phase. The Triplets model [26] contains sequential first-order phase transitions that separate a density-dependent relativistic mean-field model with nucleons and hyperons phase from a two-flavor quark

color-superconducting phase (2SC) and a quark color-flavor-locked phase. The RDF model introduced in Ref. [118], generalized to include density-dependent vector and diquark couplings in Ref. [116], contains a first-order transition that separates a hadronic regime [123] from a 2SC phase. First-order phase transitions can also be constructed phenomenologically to separate phases from different descriptions using n_B^* and Δn_B as a variable parameter to tune the transition density and the gap in baryon number density between the two phases (see, e.g., Refs. [116,118,128–131]). Such transitions are straightforward to model with a mGP assuming a Maxwell construction by replacing a portion of the baseline c_s^2 with a segment for which $c_s^2 = 0$.

III. CHOICE OF PRIORS

Now that we have described how EoS are created from the benchmark GP and the mGP models for $c_s^2(\log_{10} p)$, we will specify the relevant prior distributions. Generally, what constitutes an appropriate prior will depend on the parameters being estimated in a Bayesian analysis. The implicit assumptions we make by modeling the EoS from a nonparametric framework are that (i) the speed of sound at each sampled value of pressure, $c_s^2(p_i)$, is an effective parameter, and (ii) the method and the hyperparameters we choose for generating $c_s^2(p_i)$ dictate both the prior distribution and the correlations across the effective parameter space.

The prior distribution that we choose is, therefore, a statement on our prior beliefs of the allowed values of $c_s^2(p_i)$. When dealing with effective parameters of this type, there are two important aspects to consider. On the one hand, we must model a diverse set of functional forms of $c_s^2(p_i)$ to span a sufficiently large sample of its function space. On the other hand, we must also ensure that this diverse set of functional forms leads to $c_s^2(p_i)$ that are statistically consistent with astronomical observations of neutron stars, i.e., that, to the best of our knowledge, our prior offers a reasonable description of neutron stars.

A technical point about modeling $c_s^2(\log_{10} p)$ using GPs is the following. In the previous section, we have defined crossovers and phase transitions of different orders based on the behavior of c_s^2 as a function of p or n at the transition point (in the case of phase transitions) or in the regime where the transition happens (in the case of crossovers). However, the GP kernel is the same before and after the transition density. That means that the correlation length will be the same across all phases of matter (though we do not explicitly model different phases). This assumption is contrary to what we expect to happen in nature and to what is predicted by microscopic models, which is that different phases of matter will lead to different correlations between thermodynamic variables due to changing degrees of freedom and interactions. Even so, in the context of a Bayesian analysis, as long as the prior is sufficiently flexible to

capture the behavior of the true EoS, assuming a stationary kernel, i.e., one that does not change as a function of the independent thermodynamic variable in the GP, there will be support for the true EoS in the posterior. In future work, one could explore a piecewise or density-dependent kernel. We also recommend Refs. [84,87] for a discussion on GP kernels.

In the next two subsections, we will discuss in detail the priors on the GP and mGP hyperparameters and how we ensure the sample size is large enough and consistent with reliable astronomical observations.

A. Priors on GP and mGP hyperparameters

Let us first discuss the priors that we choose on the benchmark GP hyperparameters. These hyperparameters correspond to the correlation length (ℓ) and correlation strength (σ) at each $c_s^2(p_i)$ and the slope (m) of the mean function $\mu_i(\log_{10} p_i)$. For ℓ and σ , we choose delta-function priors that fix these parameters to unity. For m , we choose two equal-probability delta-function priors, one peaked at $m = 1.6$ and one at $m = 2$, such that 50% of the time the benchmark GP corresponds to a hard GP and 50% of the time it corresponds to a soft GP.

We will now discuss the priors on the mGP hyperparameters. As noted in Sec. II B, the mGP model introduces modifications on top of a baseline, which is modeled through the benchmark GP. The introduction of one spike is controlled by four hyperparameters: a true or false switch q_{sp} , a spike magnitude sp , a spike direction \widehat{sp} , and a spike location p_{sp} , such that $c_s^2(p_{sp}) = sp$ if a spike is present. Similarly, the introduction of one plateau is controlled by five hyperparameters: a true or false switch q_{pl} , a plateau magnitude pl , a plateau direction \widehat{pl} , a plateau width in pressure Δp_{pl} , and a plateau starting location p_{pl} , such that $c_s^2(p) = pl$ for $p \in [p_{pl}, p_{pl} + \Delta p_{pl}]$ if a plateau is present.

In this work, we will consider the introduction of one spike, one spike and one plateau, and two plateaus. This implies that every modification is controlled by a choice of the hyperparameter vector $\vec{h} = \vec{h}_1 \cup \vec{h}_2$, where $\vec{h}_1 = \{q_{sp}, q_{pl1}, q_{pl2}, \widehat{sp}, \widehat{pl}_1, \widehat{pl}_2\}$ determines whether spikes and plateaus are present and $\vec{h}_2 = \{sp, p_{sp}, pl_1, \Delta p_{pl1}, p_{pl1}, pl_2, \Delta p_{pl2}, p_{pl2}\}$ determines the properties of these modifications.

The hyperparameters in \vec{h}_2 are dependent on the choices made for \vec{h}_1 , so we first focus on \vec{h}_1 . The switch hyperparameters $\{q_{sp}, q_{pl1}, q_{pl2}\}$ determine whether a feature is present, and thus they can only be 0 or 1. Similarly, the unit vector hyperparameters $\{\widehat{sp}, \widehat{pl}_1, \widehat{pl}_2\}$ indicate the direction of a spike or a plateau (i.e., whether the modification increases or decreases the speed of sound with respect to the baseline) and can only take values of ± 1 . We consider the following seven configurations:

- (i) $\vec{h}_1 = \{0, 0, 0, \widehat{sp}, \widehat{pl}_1, \widehat{pl}_2\}$. No modification is introduced, and the mGP reduces to the benchmark model.
- (ii) $\vec{h}_1 = \{1, 0, 0, \pm 1, \widehat{pl}_1, \widehat{pl}_2\}$. A spike is introduced that goes either above or below the baseline.
- (iii) $\vec{h}_1 = \{1, 1, 0, \pm 1, \mp 1, \widehat{pl}_2\}$. A spike is introduced that goes either above or below the baseline, and a plateau is introduced, which goes in the direction opposite to the spike.
- (iv) $\vec{h}_1 = \{0, 1, 1, \widehat{sp}, \pm 1, \pm 1\}$. Two plateaus with different magnitude are introduced, with both plateaus being allowed to go above or below the baseline.

We assign equal prior probability to each of these four options, implying that 25% of our samples are from the benchmark GPs and the remaining 75% come from the mGPs (i.e., out of the total number of samples, 25% contain a single spike, 25% contain a spike and a plateau, and 25% contain a double plateau.)

The remaining hyperparameters \vec{h}_2 have specific allowed ranges, which depend on which of the above four options is drawn. In the single spike case ($\vec{h}_1 = \{1, 0, 0, \pm 1, \widehat{pl}_1, \widehat{pl}_2\}$), we must choose the height of the spike sp and its location p_{sp} , such that at the spike $c_s^2(p_{sp}) = sp$. For the location of the spike, p_{sp} , we choose a flat prior with edges at $p(n_B = 1.1n_{\text{sat}})$ and $p = 10^{37} \text{ erg cm}^{-3}$. For the height parameter, sp , we choose different priors depending on whether the spike goes above or below the baseline. If the spike is up, then we choose a flat prior with edges at $sp = 1.1c_{s,\text{benchmark}}^2(p_{sp})$ and $sp = 1$. If the spike is down, then we choose a flat prior with edges $sp = 0$ and $sp = 0.9c_{s,\text{benchmark}}^2(p_{sp})$. This choice of prior guarantees that the modified speed of sound squared is never negative, never exceeds unity, and always introduces at least a 10% modification.

In the spike plus a plateau case ($\vec{h}_1 = \{1, 1, 0, \pm 1, \mp 1, \widehat{pl}_2\}$), we must first choose the properties of the plateau to guarantee that there are no spikes within the plateaus. That is, because the plateau has a width Δp_{pl1} , when sampling the location of the spike, p_{sp} , we must not sample within $[p_{pl1}, p_{pl1} + \Delta p_{pl1}]$. Given the above, for the plateau width, Δp_{pl1} , we sample on $\log_{10} \Delta p_{pl1}$ from a flat prior on the interval $[0.12, 1.2]$, with pressure in units of erg cm^{-3} . For the plateau location, p_{pl1} , we use a flat prior with edges $p(n_B = 1.1n_{\text{sat}})$ and $p = (10^{37} \text{ erg cm}^{-3} - \Delta p_{pl1})$ to ensure the entire plateau falls within the allowed pressure range. Thus, the range of the prior for the location of the spike, p_{sp} , must be modified (from the case when there is no plateau) to $[p(n_B = 1.1n_{\text{sat}}), p_{pl1}] \cup (p_{pl1} + \Delta p_{pl1}, 10^{37} \text{ erg cm}^{-3}]$. The prior on the plateau magnitude, pl_1 , is chosen in the same way as the prior on the spike magnitude sp (see the paragraph above), but with one modification: instead of setting the edge at 0.9 or 1.1 of $c_{s,\text{benchmark}}^2(p_{sp})$, we use 0.9 or 1.1 of

$c_{s,\text{benchmark}}^2(p_{\text{pl1}})$. With this adjustment, the plateau always introduces at least a 10% modification to the benchmark from the starting point of the plateau, p_{pl1} .

In the case with two plateaus ($\hat{h}_1 = \{0, 1, 1, \hat{sp}, \pm 1, \pm 1\}$), we follow the same procedure as above, but with the following modifications. After drawing a location and width for the first plateau, we must ensure the second one is distinct (i.e., nonoverlapping), and thus it must be placed to the left or to the right of the first plateau. We enforce this constraint by first drawing Δp_{pl2} from the same flat prior as that used for Δp_{pl1} . We then remove Δp_{pl2} from the right side of the intervals $[p(n_B = 1.1n_{\text{sat}}), p_{\text{pl1}})$ and $(p_{\text{pl1}} + \Delta p_{\text{pl1}}, 10^{37} \text{ erg cm}^{-3}]$ and draw p_{pl2} from a flat prior in the interval $[p(n_B = 1.1n_{\text{sat}}), p_{\text{pl1}} - \Delta p_{\text{pl2}}) \cup (p_{\text{pl1}} + \Delta p_{\text{pl1}}, 10^{37} \text{ erg cm}^{-3} - \Delta p_{\text{pl2}}]$. This procedure is computationally efficient, and it guarantees the two plateaus do not overlap.

In Fig. 3, we show examples from each of the three groups within the mGP framework: a spike modification ($\hat{sp} = \pm$), a spike and a plateau modification [$(\hat{pl}_1, \hat{sp}) = (\pm, \mp)$], and a two-plateau modification [$(\hat{pl}_1, \hat{pl}_2) = (\pm, \pm)$]. In the top panel, the samples are represented by c_s^2 as a function of $\log_{10} P$, where p is in units of erg cm^{-3} , exactly as they were generated by the mGP. We can see that the mGP framework succeeds in introducing multiscale correlations to the speed of sound functional form. Once the samples are generated, the EoS $p(\epsilon)$ is extracted by solving the differential equation $dp/d\epsilon = c_s^2(\log_{10} p)$. In the middle panel, the same samples of c_s^2 are shown but now as a function of n_B/n_{sat} . In this panel, we also calculate the maximal central density for a stable, non-rotating star and denote it with a circle.

From these two panels, we can make several observations. First, note that in the middle panel the structure in c_s^2 is more condensed at low densities and more spread out at large densities, relative to the structure in the top panel. This is because the pressure increases more rapidly as a function of density in the outer layers of the star, where the densities are low. As a consequence, even the smallest structure in c_s^2 at low densities introduces structure over a large range of pressures. In contrast, at higher densities, the pressure increases slowly with respect to density. Therefore, dramatic features in c_s^2 at high densities translate into structure that arises over a small range in pressures.

From this comparison, we also arrive at an important conclusion regarding the optimal variable to sample over when introducing modifications. As shown in the middle panel, most stable, nonrotating stars (shown with darker lines) will reach at most $n_B \sim 7n_{\text{sat}}$. The density regime between $1 - 7n_{\text{sat}}$ is precisely where structure in c_s^2 is spread out over a larger interval in p . This implies that sampling over p will allow us to resolve this structure better than if we sampled over n_B , for a finite resolution. This is

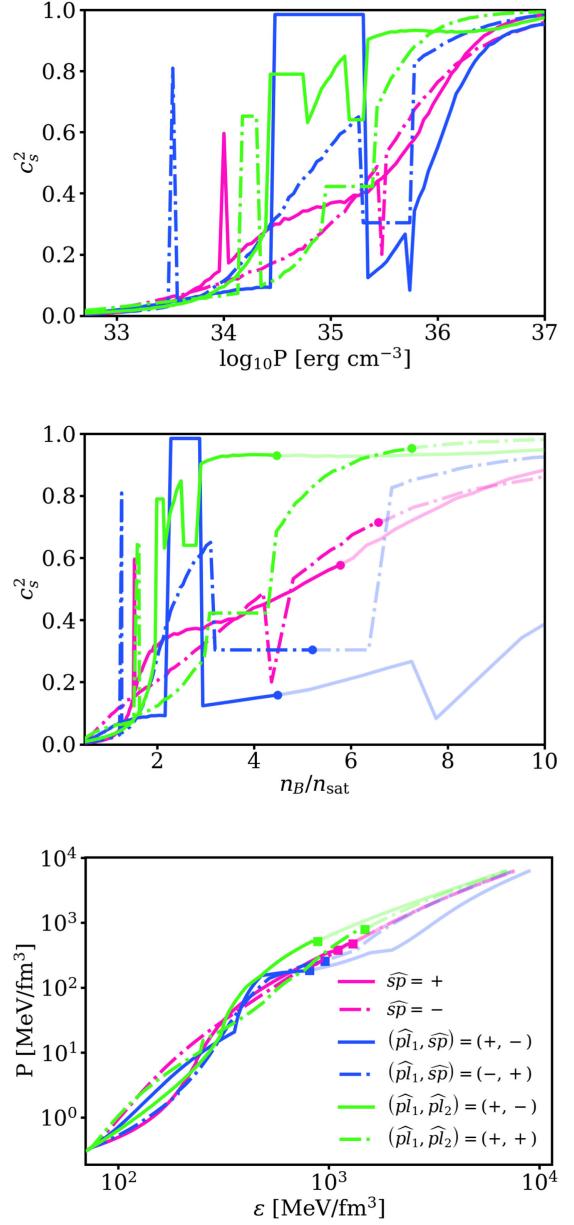


FIG. 3. Top to bottom: the speed of sound squared in units of c^2 as a function of the pressure in units of erg cm^{-3} , the speed of sound squared in units of c^2 as a function of the baryon number density in units of n_{sat} , the pressure as a function of energy density, both in units of MeV fm^{-3} , for a representative set of samples generated using the mGP framework. The circles in the middle panel represent the maximal central density predicted for a stable, nonrotating star. The squares in the bottom panel are the corresponding maximal values for the pressure and energy density for a stable, nonrotating star. Samples that contain a spike are shown in pink, samples that contain a spike and a plateau are shown in blue, and samples that contain two plateaus are shown in green. Solid and dot-dashed lines differentiate between different features for samples in the same category, as indicated by the switch parameters \hat{sp} , \hat{pl}_1 , and \hat{pl}_2 , which are defined in Sec. III C. The mGP framework produces a diverse set of EoS which contain multiscale correlations across densities at a low computational cost.

the choice, i.e., to sample in pressure, that we will make henceforth in this paper. We also note that n_B^{\max} can only be determined once the modifications have been sampled. During the sampling procedure, we reject mGP EoS for which modifications are present only above n_B^{\max} . In cases where two modifications are sampled (e.g., a spike and a plateau), we require that at least one modification is present below n_B^{\max} (an example is shown in the middle panel of Fig. 3 in solid blue).

The bottom panel of Fig. 3 displays the EoS samples on the $\epsilon - P$ plane, with squares denoting the maximal central pressure and energy density for each EoS. From this panel, it is clear that the modifications introduced in c_s^2 are not causing the EoS to significantly deviate from each other. In fact, we see that large changes in the speed of sound translate into rather small changes in the EoS, leading to clustering around a region on the $\epsilon - p$ plane.

We will now discuss in detail the role of spike modifications. Because c_s^2 is the differential of p with respect to ϵ , changes in $c_s^2(p)$ that span a short range in p may not result in visible features on the $p(\epsilon)$ plane, even if the change in the value of c_s^2 at a given point in p is significant. An example of this behavior is seen in the blue-dashed functional form in Fig. 3. The first spike is a significant deviation from the baseline value of c_s^2 as a function of both $\log_{10} p$ and n_B , but this modification has no visible effect on the functional form on the $\epsilon - p$ plane. The impact of a spike modification introduced in $c_s^2(p)$ on the corresponding EoS in $p(\epsilon)$ depends on the value of p and the value of c_s^2 where the spike is introduced with respect to the original benchmark value of c_s^2 . We discuss the interplay between these factors in the Appendix. In summary, we find that a spike modification can lead to changes in ϵ as large as $\mathcal{O}(100)$ MeV/fm³. Therefore, what we call a spike modification may be pointlike on the $\log_{10}(p) - c_s^2$ plane, but it can translate into a large modification to the EoS on the $\epsilon - p$ plane. However, it can also result in no visible change at all to the EoS on the $\epsilon - p$ plane.

A natural question at this stage is how many samples that have a spike modification are *a priori* indistinguishable from the original benchmark EoS, since these samples would make the benchmark GP and the spike-type mGP priors equivalent. Studying the relationship between observations and the scale of changes in $p(\epsilon)$ is a central component of this work. That is because how much an observable, and thus the likelihood, changes in response to a change in $p(\epsilon)$ depends on the observable type, the precision of the measurement, the individual EoS, and at what densities the modifications are introduced. In this work, we ask if current observations are sensitive to features in the EoS corresponding to state-of-the-art predictions from nuclear physics models with beyond neutron-proton-electron degrees of freedom. These features include changes in $p(\epsilon)$ with respect to a benchmark model spanning a broad range of scales, which we choose to categorize as spikes and

plateaus. To answer this question, we combine different observation types, while accounting for their uncertainties consistently and marginalizing over the uncertainties in the EoS. We then check if there is a cutoff scale that emerges from the data for the onset and/or the range over which changes in the degrees of freedom can occur (i.e., if certain features are favored/disfavored). Because our approach is data driven, our priority is to sample broad range of features and let the data determine if features can be distinguished from each other, i.e., favored/disfavored.

One could, instead, pursue a model-based approach and check how different modifications of $c_s^2(\log_{10} p)$ or $c_s^2(n_B)$ translate into changes in the $p(\epsilon)$ EoS and, consequently, the mass-radius or tidal deformability-compactness curves. The downside of taking this approach is that, while one could consider observational uncertainties, i.e., one can argue that a modification is only relevant if it produces a deviation with respect to the benchmark that would be detectable taking into account current experimental uncertainties, it does not make use of all available observations and their uncertainties. Also, because the effect of a modification on the $\epsilon - p$ plane depends on all the factors mentioned above, systematically checking all the possibilities would be a computational and analytical challenge. Therefore, the approach taken in this work is appropriate because it incorporates all the available experimental, theoretical, and observational information we have on beta-equilibrated, cold nuclear matter and allows for adequate sampling of the features predicted by nuclear physics models in our priors.

We stress here that Fig. 3 only shows six representative samples of c_s^2 and EoS out of the nearly one million samples that we study in this paper.

B. Astronomical observables

We have specified a prior for the EoS, a quantity that we cannot directly measure. Thus, we need to translate the information carried by the EoS into astronomical observables if we wish to infer properties of the former. For each EoS, we can calculate the mass-radius ($M - R$), moment of inertia-mass ($I - M$), quadrupole moment-mass ($Q - M$), and tidal deformability-mass curves ($\Lambda - M$), which encode the properties of neutron stars of different central densities. These properties can then be compared to astronomical observations to determine the validity of the EoS. In this subsection, we will discuss briefly how these quantities are obtained, following mostly Ref. [132]. In the next subsection, we will explain how we use astronomical observations of these quantities to inform our prior sample size.

Millisecond pulsars rotate slowly compared to their mass-shedding limit, and any radius corrections resulting from rotation are significantly below measurement precision for the NICER mission [56]. Therefore, for calculating the mass-radius curve, we can approximate millisecond pulsars

as nonrotating, isolated objects. The condition that a stable nonrotating star must be in hydrostatic equilibrium yields the well-known Tolman-Oppenheimer-Volkoff (TOV) equation. For each EoS, given a central pressure p_c , the TOV equation will describe how the mass and pressure vary with the radial coordinate up to a limiting value, R_* , where $p(R_*) \approx 0$. The mass of the star is given by $M = M(R_*)$, and the stellar radius is $R = R_*$. By repeating this process for multiple values of p_c , we get the $M - R$ curve for a given EoS.

At some value of central pressure, the $M - R$ sequence will become unstable, which occurs when $dM/dp_c < 0$. This value of the central pressure marks the end of what we refer to as the (stable) neutron star EoS and establishes a maximal allowed value for the central pressure, p_c^{\max} , for astrophysically realistic stars. Note that the zero-temperature QCD EoS continues beyond p_c^{\max} , but we do not expect to observe stable, isolated, nonrotating or slowly rotating pulsars with central pressures exceeding p_c^{\max} , assuming the EoS used in the calculation is correct. We also note that the $M - R$ sequence can have multiple stable branches separated by unstable regimes where $dM/dp_c < 0$. When that is the case, p_c^{\max} is still the largest density corresponding to the end point of a stable branch, but it might not correspond to the central density of a maximally massive star for a given EoS.

Now, let us consider an isolated star that rotates with dimensionless angular velocity Ω , where we have normalized the stellar angular velocity by the mass-shedding limit. We assume Ω is small enough that the solutions to the Einstein equations can be expanded in powers of this dimensionless angular frequency. Note that to $\mathcal{O}(\Omega^0)$ the (t, t) and (r, r) components of the Einstein equations, in conjunction with the stress-energy conservation equation, yield the TOV and continuity equations. At $\mathcal{O}(\Omega)$, the Einstein equations only modify the gravitomagnetic sector of the metric, whose exterior behavior is characterized by the moment of inertia I . At $\mathcal{O}(\Omega^2)$, we obtain a correction to the total mass and to the mass distribution within the star, which now acquires an oblate spheroidal shape due to rotational motion. The latter is controlled by the quadrupole moment Q . Both I and Q vary with the central density of the star, such that we can obtain solutions for a range of central pressures and relate them to the $M - R$ curve to obtain the $I - M$ and $Q - M$ curves.

Finally, we will consider a nonrotating star in the presence of a companion compact object, which causes the star to tidally deform. We can study the redistribution of mass due to the external perturbation through a multipolar expansion. The deformation at leading order in perturbation theory is dominated by the quadrupole moment. The quadrupolar deformation is controlled by the Love number, or its dimensionless version, the tidal deformability Λ . The tidal deformability can be calculated by solving the

linearized Einstein equations combined with continuity and differentiability arguments. Once again, the exact solution for Λ requires an EoS and is dependent on the central pressure of the star [133]. The calculation can be repeated for a range of central pressures to obtain the $\Lambda - M$ curve.

C. Prior sample size

The priors on neutron star observables, such as their mass, radius, and tidal deformability are determined by the prior on the EoS. But how do we know that we have chosen a good EoS prior? For each measurement available, we want enough samples in the prior that offer a reasonable match to observations as predictions. Because of the functional complexity allowed by both the GP and mGP frameworks, we expect that most EoS generated will not meet basic requirements based on neutron star observations. With that in mind, we use three metrics to gauge the how well samples in the prior describe astronomical measurements. First, we check that $M_{\max} \geq 1.8M_{\odot}$ based on the observation of three high-mass pulsars [52–54]. Second, based on the inference of the tidal deformability of a $1.4M_{\odot}$ star, $\Lambda_{1.4}$, from GW170817 [134], we check that $10 \leq \Lambda_{1.4} \leq 2000$. Third, based on NICER's inference of the mass-radius posterior for PSR J0030 + 0451 [55,57], we check that $8.0 \leq R_{1.4} \leq 16.0$ km. We note that these bounds are far outside the 90% credible region for the most constraining estimates of the maximum mass [135], radius of a $1.4M_{\odot}$ star [56], and tidal deformability of a $1.4M_{\odot}$ star [134].

For every sample EoS that we draw from our GP or mGP framework, we keep tally of whether the three checks specified above are passed or not. We emphasize that these metrics are *not* used to cut the sample size or to modify our priors in any way. We simply track how many samples pass these checks to ensure that we have enough strong candidate EoS in our prior sample. In particular, we continue drawing samples until we have obtained a subset of at least 100,000 candidate EoS that pass all three checks. This requires that we sample about 1,000,000 times from the benchmark GP and mGP frameworks.

Another benefit of checking our priors in terms of astronomical observables is that we can assign zero likelihood to EoS that fall outside the intervals we defined above. This is because those sample EoS are already in significant conflict with the observations discussed above, and thus their likelihood will be very close to zero. We can justify this approach as follows. Consider an observable Y at the value y_k predicted by EoS k , which we will parametrize in terms of a vector $\vec{\phi}_k$ [see Eq. (2), where here the vector symbol denotes the 100 values of ϕ that we sample at each point in pressure]. Let us also consider a set of N total number of EoS, such that k is between 1 and N . Then, the conditional probability of y_k given EoS k is

$$P(y_k) = \frac{q(\vec{\phi}_k)\mathcal{L}(\vec{\phi}_k)}{\sum_i^N q(\vec{\phi}_i)\mathcal{L}(\vec{\phi}_i)}, \quad (11)$$

where $q(\vec{\phi}_k)$ is the prior probability assigned to EoS k and $\mathcal{L}(\vec{\phi}_k)$ is the likelihood of the data given EoS k . Let us now order the values of y_k from smallest to largest, such that

$$P(y_1 \leq y \leq y_2) = \sum_k P(y_k), \quad y_k \in [y_1, y_2] \quad (12)$$

defines the credible region delimited by y_1 and y_2 . For most observables, only a limited domain in y_k will have nonzero likelihood. That is,

$$P(y_{\text{low}} \leq y \leq y_{\text{high}}) = \sum_k P(y_k) \approx 1, \quad y_k \in [y_{\text{low}}, y_{\text{high}}], \quad (13)$$

where $P(y_k) \approx 0$ for y_k outside the interval $[y_{\text{low}}, y_{\text{high}}]$. For N_{obs} different observables (e.g., maximum mass, radius at $1.4M_{\odot}$, or dimensionless tidal deformability at $1.4M_{\odot}$), we then simply have an N_{obs} -dimensional region $[y_{\text{low}}^1, y_{\text{high}}^1] \cup [y_{\text{low}}^2, y_{\text{high}}^2] \cup \dots \cup [y_{\text{low}}^{N_{\text{obs}}}, y_{\text{high}}^{N_{\text{obs}}}]$ outside of which we expect $P(y_k) \approx 0$. In our case, we consider three observables, so the region outside of which the posterior is approximately zero is $[M_{\text{max,low}} = 1.8M_{\odot}, M_{\text{max,high}} = \infty) \cup [R_{1.4,\text{low}} = 8.0 \text{ km}, R_{1.4,\text{high}} = 16.0 \text{ km}] \cup [\Lambda_{1.4,\text{low}} = 10, \Lambda_{1.4,\text{high}} = 2000]$.

We can implement the condition that the posterior is zero outside of the above region as follows. First, we divide our prior sample into two subsets, one for which each EoS meets all the requirements specified above (Φ_{\vee}), where

$$\Phi_{\vee} = \{\vec{\phi}_k : M_{\text{max,k}} \geq 1.8M_{\odot} \wedge R_{1.4,k} \in [8.0 \text{ km}, 16.0 \text{ km}] \wedge \Lambda_{1.4,k} \in [10, 2000]\}, \quad (14)$$

and one for which all EoS fail at least one of the checks (Φ_{\times}), such that

$$\Phi = \Phi_{\vee} \cup \Phi_{\times}. \quad (15)$$

With this in hand, we now define the likelihood of the data given EoS k to be $\mathcal{L}(\vec{\phi}_k)_{\vee} = \mathcal{L}(\vec{\phi}_k)$ if $\vec{\phi}_k \in \Phi_{\vee}$, and we define $\mathcal{L}_{\times}(\vec{\phi}_k) = 0$ if $\vec{\phi}_k \in \Phi_{\times}$. We emphasize again that the prior distribution remains unchanged, so this procedure is in no way equivalent to performing cuts on the prior.

Our goal is to generate enough samples to ensure that Φ_{\vee} contains at least $\sim 100,000$ EoS. Using the M_{max} , $R_{1.4}$, and $\Lambda_{1.4}$ checks as a guide, we generate 900,000 EoS. Out of this total sample, 104,594 EoS passed the checks and therefore contribute non-negligibly to the posterior distribution of the observables discussed later in Sec. V. Note that the number of samples in Φ_{\vee} is roughly 10% of the total number of samples generated. Based on this result, we argue that studies using nonparametric methods, or any

method that allows for a vast functional space, should implement similar checks, or at least verify the robustness of results for different prior sample sizes.

IV. STATISTICAL METHODS

There should be a unique EoS that correctly describes all neutron stars in the Universe. However, honing in on this exact EoS would only be possible with infinitely precise observations. A more common and realistic approach is to obtain posteriors that describe the probability of a given EoS by comparing its predictions against data. Using an ensemble of theoretical models for the EoS, each with a corresponding posterior probability, we can extract credible regions for the EoS that occurs in nature. This method requires us to first state our prior beliefs about the EoS, which then get updated as we gain knowledge of the EoS through data.

We have introduced in Sec. II two frameworks for generating theoretical models for the EoS: benchmark GPs and mGPs. Those frameworks reflect two different beliefs about the EoS. The GP assumes the EoS displays long-range correlations in pressure, resulting in smooth functional forms for $c_s^2(p)$. This belief is compatible with nuclear physics simulations for hadronic models, or models that display a smooth crossover into an exotic phase, where the change in the degrees of freedom happens over a wide range in density. On the other hand, the mGP framework assumes the EoS contains nontrivial degrees of freedom or interactions that lead to sudden changes in $c_s^2(p)$ in the form of kinks, spikes, and plateaus. These features are predicted by many state-of-the-art nuclear physics simulations with exotic degrees of freedom. With these two distinct prior beliefs in mind, our goal is to assess if one framework is better at accounting for observations than the other.

We attempt to answer this question using a fully Bayesian approach. In Sec. III, we detailed how we generate a prior distribution using the benchmark GP and the mGP as theoretical frameworks. Now, we need to discuss how we calculate posterior distributions by incorporating constraints on the EoS of neutron stars from astronomical observations, controlled terrestrial experiments, and perturbative QCD calculations,⁴ and how we quantify each framework's ability to describe observations.

We begin this section with a Bayesian “primer,” where we review a generic approach for implementing our knowledge about the EoS into a posterior distribution and how we can determine whether observations favor one of the EoS frameworks over the other using the Bayes factor. Obtaining posterior distributions requires specific choices and assumptions for which observations and associated

⁴For a detailed discussion on current constraints on the QCD EoS across different regimes, see Ref. [119].

likelihood factors are used. Those are discussed in a dedicated likelihood subsection. Similarly, we devote a separate subsection to explaining how the model evidence for the benchmark GP and the mGP are determined—a requirement for calculating the Bayes factor.

A. Bayesian primer

Consider an EoS k which is represented by a set of values sampled from either the benchmark GP or the mGP, $\vec{\phi}_k$. Bayes's theorem states that the posterior probability of EoS k is proportional to the product of a prior term and the likelihood of the data given $\vec{\phi}_k$,

$$P_{\text{EoS}}(\vec{\phi}_k) \propto q(\vec{\phi}_k) \mathcal{L}(\vec{\phi}_k), \quad (16)$$

where we recall that $q(\vec{\phi}_k)$ is the prior probability distribution encoding our prior beliefs about how likely EoS k is to occur and we recall that the likelihood term $\mathcal{L}(\vec{\phi}_k)$ reflects how well predictions from EoS k match observed properties.

In this paper, we assume that all observations are independent of each other, which means the likelihood of a set of observations (i, j) for EoS model k can be written as

$$\mathcal{L}(\vec{\phi}_k) = \prod_i \left[\prod_{j=1}^{j(i)} \mathcal{L}_k(i, j) \right], \quad (17)$$

where i is a type of measurement (e.g., mass, radius) and j is an independent measurement of type i (e.g., two independent measurements of the mass of one object). For each of the measurements, we must make a choice for how it will be incorporated into the analysis via a likelihood function, $\mathcal{L}_k(i, j)$. We will discuss our choice of likelihood functions and specify which measurements we include in our analysis in the next subsection.

Let us now instead return to Eq. (16). The normalizing factor that would make Eq. (16) an equality is called the model evidence. The model evidence assesses the ability of a set of prior beliefs to account for observations. In the case of nonparametric EoS, the evidence can be defined as

$$\mathcal{E}_m = \int_{\Phi_m} \mathcal{L}(\vec{\phi}_k) q(\vec{\phi}_k) dk, \quad (18)$$

where Φ_m is the set of samples in the prior that were generated using a specific theoretical framework m . The goal is often to have competing frameworks such that we can compute the model evidence for each one and then take the ratio between them. This ratio between model evidences is known as the Bayes factor,

$$K = \frac{\int_{\Phi_{m1}} \mathcal{L}(\vec{\phi}_{k1}) q(\vec{\phi}_{k1}) dk_1}{\int_{\Phi_{m2}} \mathcal{L}(\vec{\phi}_{k2}) q(\vec{\phi}_{k2}) dk_2}, \quad (19)$$

where $m1$ and $m2$ indicate distinct theoretical frameworks with different equations of state in their samples (indexed here by k_1 and k_2). When the Bayes factor deviates significantly from unity, it indicates that the data prefer one model and prior over the other.

B. Likelihood

As stated in Eq. (17), we assume that all measurements we take into account are independent and that systematic errors can be neglected such that the total likelihood is a product of individual likelihood factors for each measurement. In particular, we will consider estimates of the nuclear symmetry energy, the three highest reliably measured pulsar masses, two NICER simultaneous mass and radius measurements, and tidal deformability estimates from two gravitational-wave events. Additionally, we will incorporate a perturbative QCD weight [67], which accounts for the behavior of the EoS at very large ($\sim 40n_{\text{sat}}$) densities from pQCD calculations. In summary, Eq. (17) can then be written as

$$\mathcal{L}(\vec{\phi}_k) = \mathcal{L}_S(\vec{\phi}_k) \mathcal{L}_{M\text{max}}(\vec{\phi}_k) \mathcal{L}_{M-R}(\vec{\phi}_k) \mathcal{L}_\Lambda(\vec{\phi}_k) w_{\text{pQCD}}(\vec{\phi}_k), \quad (20)$$

where S denotes the likelihood factor associated with symmetry energy measurements, M_{max} denotes that associated with high-mass pulsar measurements, $M - R$ denotes that associated with simultaneous mass-radius measurements, and Λ denotes that associated with tidal deformability measurements. We represent input from pQCD not as an additional likelihood factor but as a weight, w_{pQCD} . We make this choice because the uncertainty in the pQCD input stems from its poorly constrained regime of applicability and uncertainty around the missing higher-order term error in truncated results, in contrast to traditional measurements with quantifiable statistical uncertainties that can be consistently included in a Bayesian framework. Reference [136] demonstrated that Bayesian inference techniques can be used to quantify the uncertainty in the pQCD input. However, more work is needed to determine if these results are sensitive to the choice of uncertainty estimation technique. Our approach for incorporating observational and experimental constraints on the EoS is based on Refs. [55, 56, 68], while the use of pQCD input is based on Ref. [67]. We review and discuss the most important aspects of these approaches as they pertain to our analysis below and refer the reader to the corresponding original works for further detail.

1. Symmetry energy

In terrestrial experiments, it is possible to probe the $T \rightarrow 0$ limit of dense nuclear matter with low-energy collisions of heavy ions [119]. However, the nuclei used in these experiments have a charge fraction (the ratio of proton number to baryon number or, in other words, electric charge density n_Q over baryon density n_B) of $Y_Q = n_Q/n_B \sim 0.4\text{--}0.5$. A value of $Y_Q = 0.5$ is known as SNM because there is an equal number of protons and neutrons in the system. On the contrary, neutron stars are primarily neutron rich, with $Y_Q \sim 0.001\text{--}0.2$, thus probing the asymmetric nuclear matter regime. PNM is the limit where $Y_Q = 0$.

The densities probed in these low-energy heavy-ion experiments are at or near n_{sat} , and in that regime, experiments and χ EFT calculations can extract properties that are relevant to the EoS. A nucleus is composed of Z protons and $A - Z$ neutrons, where A is the total number of protons and neutrons in the nucleus. The mass of the nucleus m_A that contains A nucleons is always less than the masses of the individual protons, m_p , and neutrons, m_n , summed together, i.e.,

$$m_A < Zm_p + (A - Z)m_n, \quad (21)$$

because a finite amount of energy is released in the formation of a nucleus. That difference in the rest mass energy per nucleon is known as the binding energy and is defined as (we remind the reader that the speed of light is $c = 1$ in this work)

$$B = \frac{1}{A} [m_A - (Zm_p + (A - Z)m_n)]. \quad (22)$$

It is often assumed that the mass of the proton and neutron are identical since their masses differ by just over 1 MeV. Setting the neutron and proton mass to be the same and calling this the nucleon mass for simplicity, $m_p = m_n = m_N$, the above simplifies to

$$B = \frac{m_A}{A} - m_N \equiv \frac{E}{A}, \quad (23)$$

where E/A is the nucleonic energy per particle. As defined here, the binding energy does not depend explicitly on the mass of the nucleon because the m_N dependence in the first term of the above equation cancels the second term exactly.

The binding energy for SNM at n_{sat} is estimated to be $B \sim -16$ MeV from previous global analyses⁵ that extracted the volume term of the liquid drop model from a large sample of nuclei, which reported values of $B = -15.77$ [137] and $B = -16.24$ MeV [138]. One can also use χ EFT tuned to a large number of experimentally

measured nuclei wherein one obtains $B = -15.86 \pm 0.57$ MeV [62]. However, in this work, we assume $B = -16$ MeV is exact.

The next quantity that can be measured from nuclear experiments is known as the symmetry energy, which we denote as S . At n_{sat} , the symmetry energy is the difference in total energy between the SNM and PNM limits, i.e.,

$$S(n_{\text{sat}}) = \frac{1}{A} (E_{\text{PNM}} - E_{\text{SNM}}), \quad (24)$$

or, in terms of energy densities, we can write

$$S(n_{\text{sat}}) = \frac{1}{n_{\text{sat}}} (\varepsilon_{\text{PNM}} - \varepsilon_{\text{SNM}}), \quad (25)$$

where we can relate the energy densities to the total energy via

$$\frac{\varepsilon}{n_B} = \frac{E}{A} + m_N. \quad (26)$$

Neutron stars are not exactly in the limit of PNM since a small fraction of protons exists is also present. Thus, for asymmetric nuclear matter (ANM), where the system is at finite value of Y_Q , the symmetry energy can be expanded about $Y_Q = 1/2$ to obtain

$$S_2(n_B)[1 - 2Y_Q]^2 + \mathcal{O}[1 - 2Y_Q]^4 = \frac{1}{n_B} (\varepsilon_{\text{ANM}} - \varepsilon_{\text{SNM}}), \quad (27)$$

where the factor of 1/2 in the Taylor expansion is reabsorbed into the quadratic term $S_2(n_B)$. The quadratic coefficient of the Taylor expansion, $S_2(n_B)$, can then be further Taylor expanded about $n_B = n_{\text{sat}}$, but in this paper, we will retain only the leading-order term in this expansion and set $S(n_{\text{sat}}) = S_2(n_{\text{sat}})$ in Eq. (26) (see Ref. [60] for a derivation and further details).

We can relate the symmetry energy to the binding energy by substituting in Eq. (26) into Eq. (27) at $n_B = n_{\text{sat}}$ for the $\varepsilon_{\text{SNM}}/n_{\text{sat}}$ term to find

$$S(n_{\text{sat}})[1 - 2Y_Q]^2 \approx \frac{\varepsilon_{\text{ANM}}}{n_{\text{sat}}} - \left[\frac{E}{A} + m_N \right] \quad (28)$$

$$= \frac{\varepsilon_{\text{ANM}}}{n_{\text{sat}}} - [B + m_N]. \quad (29)$$

In an extreme extension of Taylor expansions, however, we will evaluate the above expression at very small Y_Q because state-of-the-art χ EFT models predict a value for $S(n_{\text{sat}})$ and indicate that $Y_Q \sim 0.05$ at n_{sat} for β -equilibrated, cold nuclear matter [139]. Therefore, setting $Y_Q = 0$, $B = -16$ MeV and $m_N = m_n = 939.6$ MeV in the above equation, we obtain

⁵Uncertainty quantification was not performed in these studies.

$$S(n_{\text{sat}}) \sim \frac{\epsilon_{\text{PNM}}}{n_{\text{sat}}} - 923.6 \text{ [MeV].} \quad (30)$$

Given all of the above, we assume a likelihood factor associated with the symmetry energy of the form

$$\begin{aligned} \mathcal{L}_S(\vec{\phi}_k) &= \mathcal{L}(S_0|, S_k(n_{\text{sat}})) \\ &= \frac{1}{\sqrt{2\pi(\sigma_S)^2}} \exp \left[-\frac{(S_k(n_{\text{sat}}) - S_0)^2}{2(\sigma_S)^2} \right], \end{aligned} \quad (31)$$

where $S_k(n_{\text{sat}})$ now denotes the EoS $\vec{\phi}_k$ from which the $S(n_{\text{sat}})$ is obtained using Eq. (30), assuming that $\epsilon_{\text{PNM}} = \epsilon_k(n_{\text{sat}})$, and we take the observed value of the symmetry energy at nuclear saturation density to be $S_0 = 32 \text{ MeV}$ with a standard deviation of $\sigma_S = 2 \text{ MeV}$ [59,60].

Future work could consider other available constraints [60,61] on the symmetry energy. Additionally, one could fold into the analysis the uncertainty on the value of the binding energy, the uncertainty on the value of n_{sat} itself, the systematic error introduced by neglecting deviations from exact charge fraction asymmetry (i.e., the small Y_Q that we set to zero in this derivation), and the slope of the symmetry energy. The impact of uncertainties around the binding energy, saturation density, and small deviations from PNM at n_{sat} mostly affect stars with masses below $1.0M_\odot$ [68]. A more promising direction is to include the slope of the symmetry energy, L , which would provide constraints when $n_B \neq n_{\text{sat}}$. However, the exact value of the slope of the symmetry energy is poorly constrained with current theory and experiments. We also refer the reader to Refs. [91,140], which extract from a GP prior and neutron star measurements values for S_0 , L , and the neutron skin thickness of ^{208}Pb .

2. High-mass pulsars

In principle, we could gain information about the EoS from any neutron star mass measurement. Given a fully parametrized model for neutron star birth and accretion and a prior for the neutron star EoS, we could make population predictions that can be compared to measurements. However, due to both computational and theoretical challenges, such a complete analysis is not currently feasible. Without any information about how stars form and accrete, we can focus instead on the maximum mass, since any realistic EoS must be able to support the highest reliably measured masses. As discussed in Sec. III B, in the slow rotation regime, the maximum mass is a function only of the EoS. We implement a likelihood factor for EoS k given a likelihood function for the mass of star j obtained from a radio observation of a binary pulsar, $\mathcal{L}(M_j|M(\rho_c))$, that can be written as

$$\mathcal{L}(M_j|k) = \int_{\rho_{\text{min}}}^{\rho_{\text{max}}} q(\rho_c) \mathcal{L}(M_j|M(\rho_c)) d\rho_c, \quad (32)$$

where ρ_c is the central density. The prior on the central density, $q(\rho_c)$, is calculated for each EoS from a distribution that is quadratic between the central density of a $1M_\odot$ star, which we denote ρ_{min} , and that of a maximally massive star, ρ_{max} , assuming EoS k . That is, we sample uniformly between $0 \leq x \leq 1$ for $\rho_c = \rho_{\text{min}} + x^2(\rho_{\text{max}} - \rho_{\text{min}})$. We use a quadratic prior to avoid giving greater prior weight to more massive stars, since the central density changes more rapidly near the maximum mass [56]. Outside the interval $[\rho_{\text{min}}, \rho_{\text{max}}]$, the prior probability is zero. We also assign zero prior probability for any ρ_c resulting in an unstable star, such that if two or more stable branches are connected by an unstable branch the unstable branch does not contribute to the likelihood.

We model the posterior probability distribution for the mass of an observed pulsar with a Gaussian function, namely,

$$P(M_j|M) = \frac{1}{\sqrt{2\pi(\sigma_{M_j})^2}} \exp \left[-\frac{(M - M_j)^2}{2(\sigma_{M_j})^2} \right], \quad (33)$$

where M_j is the maximum likelihood estimate for the mass of a given pulsar and σ_{M_j} is the standard deviation for that observation. The total likelihood associated with heavy pulsar mass measurements is then

$$\mathcal{L}_{M_{\text{max}}}(\vec{\phi}_k) = \prod_j \mathcal{L}(M_j|M_{\text{max},k}). \quad (34)$$

This form for the likelihood only disfavors EoS with $M_{\text{max},k}$ less than existing observations. We do not want to disfavor EoS with $M_{\text{max},k}$ higher than existing observations because observations may be biased toward lower masses for a variety of reasons unrelated to the EoS.

We incorporate in our analysis the two highest mass measurements of neutron stars in relativistic binaries, which allow for measurements of post-Keplerian parameters, such as the Shapiro delay, pericenter precession, and orbital decay due to the emission of gravitational radiation (currently the gold standard for mass measurements of neutron stars [68]). In particular, we use measurements of $M_{\text{J1614-2230}} = 1.908M_\odot$ with $\sigma_{M_{\text{J1614-2230}}} = 0.016M_\odot$ for PSR J1614-2230 [52,141] and $M_{\text{J0348+0432}} = 2.01M_\odot$ with $\sigma_{M_{\text{J0348+0432}}} = 0.04M_\odot$ for PSR J0348 + 0432 [53]. There is a third pulsar, PSR J0740 + 6620, with a measured mass of $M_{\text{J0740+6620}} = 2.07M_\odot$ and $\sigma_{M_{\text{J0740+6620}}} = 0.08M_\odot$ [54,142] that we will also include in the next subsubsection as a joint mass-radius measurement. We also note that the mass measurement for PSR J1614 – 2230 was recently updated and reported to be $M_{\text{J1614-2230}} = 1.937M_\odot$ with $\sigma_{M_{\text{J1614-2230}}} = 0.014M_\odot$ [135], but we do not expect this update to affect our results significantly. In Sec. IV B 6, we explain why we leave out of our analysis the even-higher

neutron star masses that have been suggested in the literature.

Lastly, we note that the likelihood-based approach used here (emphasized in Ref. [68]) is preferable in a Bayesian analysis compared to imposing a lower bound on the maximum mass for two reasons. First, a hard bound does not account for uncertainties in the mass measurements. As Ref. [68] illustrates, if we consider the mass estimate for PSR J0740 + 6620, $M_{\text{J0740+6620}} = 2.07 \pm 0.08 M_{\odot}$, a 1- σ maximum mass cut at $M_{\text{max}} \geq 2.04 M_{\odot}$ predicts that EoS with maximum masses of $2.04 M_{\odot}$ and $2.14 M_{\odot}$ are equally viable, when in reality, assuming Gaussian statistical uncertainties, the latter is substantially more consistent with observations. Second, a hard bound does not allow for the incorporation of information from multiple measurements. Thus, although widely practiced in the literature, imposing maximum mass cuts is statistically inconsistent and discards important information.

3. NICER

Still assuming slow rotation, the $M - R$ curve for an EoS specifies a prediction for the radius given a stellar mass which is only a function of the EoS itself. Thus, given a joint $M - R$ posterior, we can integrate over central densities and the full $M - R$ sequence predicted by EoS k to obtain the likelihood factor associated with an independent radius measurement l ,

$$\mathcal{L}(R_l|R_k) = \int_{\rho_{\min}}^{\rho_{\max}} q(\rho_c) \mathcal{L}_l(M(\rho_c), R_k(M(\rho_c))) d\rho_c, \quad (35)$$

where $\mathcal{L}_l(M(\rho_c), R_k(M(\rho_c)))$ is the likelihood of a mass $M(\rho_c)$ and a radius $R(M(\rho_c))$ given measurement l , $R_k(M(\rho_c))$ is the circumferential radius for a star with gravitational mass M given EoS k and central density ρ_c , and $q(\rho_c)$ is the prior on the central density, which is determined for each EoS in the same way as described in the previous section.

The total likelihood associated with simultaneous mass and radius measurements is then

$$\mathcal{L}_{M-R}(\vec{\phi}_k) = \prod_l \mathcal{L}(R_l|R_k). \quad (36)$$

This particular form for the likelihood is equivalent to integrating the full mass and radius likelihood over the full $M - R$ sequence predicted by an EoS. It accounts for measurement uncertainties and possible correlations between radius and mass⁶ as well as the entire $M - R$ sequence, not just an individual $R_k(M)$.

⁶Correlations between mass and radius are system dependent and may or may not be present; see Refs. [143,144] for more details.

We adopt as constraints on the radius the posteriors obtained from NICER measurements for PSR J0030-0451 [55] and PSR J0740 + 6620 [56] (again, see, respectively, Refs. [57,58] for independent analyses of these two pulsars from a separate group within the NICER collaboration).

Though other neutron star radii estimates are available, there are potentially significant systematic errors that have not been resolved [145,146]. In contrast, NICER posteriors rely on fits of rotating hot spot patterns for which studies using synthetic waveforms found no significant mass or radius bias in statistically good fits [143,144]. Lastly, we highlight that NICER posteriors for both pulsars are non-trivial shapes on the $M - R$ plane and display significant correlations between mass and radius, further emphasizing the importance of this particular approach to calculating the likelihood.

4. Gravitational waves

As discussed in Sec. III B, for each EoS k , we can calculate the $\Lambda - M$ curve, such that for a star of gravitational mass M and equatorial radius R the tidal deformability is

$$\Lambda_k = \frac{2}{3} k_2 \left(\frac{Rc}{GM} \right)^5, \quad (37)$$

where k_2 is the tidal love number, which depends intrinsically on the EoS [147,148].

In practice, it is more constraining to incorporate input from gravitational-wave observations using information from the binary tidal deformability, which can be measured to higher accuracy. In the Taylor family of post-Newtonian waveforms, given a binary neutron star system of stars with masses M_1 and $M_2 \leq M_1$ with tidal deformabilities Λ_1 and Λ_2 , the most easily measurable quantity is [149]

$$\tilde{\Lambda} = \frac{16}{13} \frac{(M_1 + 12M_2)M_1^4\Lambda_1 + (M_2 + 12M_1)M_2^4\Lambda_2}{(M_1 + M_2)^5}, \quad (38)$$

which is sometimes called the binary or effective tidal deformability. Similarly, it can be difficult to extract individual masses from gravitational-wave events, but the chirp mass, $M_{\text{ch}} = (M_1 M_2)^{3/5} / (M_1 + M_2)^{1/5}$, can be measured with high precision since it relates directly to the gravitational-wave frequency during the inspiral phase.

Assuming a gravitational-wave event n results in a full posterior in $(M_1, M_2, \tilde{\Lambda})$ space, our procedure for incorporating it is as follows. The total likelihood factor has the form

$$\mathcal{L}_\Lambda = \prod_n \mathcal{L}(\tilde{\Lambda}_n | \tilde{\Lambda}_k) = \prod_n \int dM_1 q(M_1) \int q(M_2 | M_{\text{ch},n}, M_1) \mathcal{L}_n(M_1, M_2, \tilde{\Lambda}_k) dM_{\text{ch},n}, \quad (39)$$

where $q(M_2 | M_{\text{ch}}, M_1)$ is the prior probability density for M_2 at the value of M_2 implied by M_{ch} and M_1 and the integral is over the probability distribution for M_{ch} obtained from the gravitational-wave analysis. Since there is a limited width for M_{ch} which is dependent both on the EoS and the prior for the masses, our implementation is as follows. For a binary event involving two masses $M_1 \geq M_2$, we select the central density of a $1M_\odot$ star for the lower-mass star, $\rho_{c,2} = \rho_{\min}$. Then, we calculate the value of M_1 implied by M_{ch} , which we know to the precision that we know the chirp

mass. There is a range in $\rho_{c,1}$ that, given the value of M_2 implied by $\rho_{c,2}$, leads to M_{ch} within the 68% credible region inferred for event n . That is the range we integrate over for $\rho_{c,1}$, using the same prior as before (quadratic between ρ_{\min} and ρ_{\max}). We then select a new $\rho_{c,2}$, also following the quadratic prior, and repeat the same process for $\rho_{c,1}$ that we just outlined. We continue to increase $\rho_{c,2}$ up to the density at which M_{ch} implies $M_1 = M_2$. That means we can rewrite Eq. (39) for a single event n in terms of the central densities of the two objects,

$$\mathcal{L}(\tilde{\Lambda}_n | \tilde{\Lambda}_k) = \int d\rho_{c,1} q(\rho_{c,1}) \int q(\rho_{c,2} | (M_{\text{ch},n}, M_1(\rho_{c,1}))) \mathcal{L}_n(M_1(\rho_{c,1}), M_2(\rho_{c,2}), \tilde{\Lambda}_k | M_1(\rho_{c,1}), M_2(\rho_{c,2})) d\rho_{c,2}, \quad (40)$$

where $q(\rho_{c,2} | (M_{\text{ch},n}, M_1(\rho_{c,1}))) = 0$ outside of the 68% credible region for $M_{\text{ch},n}$. We highlight that, even though M_{ch} is typically measured to high precision, it is not statistically consistent to write the integral over $\rho_{c,2}$ as a delta function. That is because the range of M_1 allowed for a given M_2 and M_{ch} depends on both the EoS and the prior for the central densities. Consequently, it will vary between individual EoS and must be calculated separately for each EoS [68].

We include in our analysis binary tidal deformability estimates from GW170817 [48–50] and GW190425 [51]. We use the publicly available posteriors over the full model parameter space [150,151] as input for a kernel density estimate of the marginalized posterior for M_1 and $\tilde{\Lambda}$. Since the combined mass in GW190425 is high enough that one of the objects might have been a black hole, we check whether for the EoS and central density under consideration the higher-mass object is a neutron star. If so, we compute the tidal deformabilities of both stars, using the same EoS, following the procedure outlined in Refs. [147,148]. However, if the EoS predicts a black hole at the central density under consideration for the higher-mass object, we set its tidal deformability to zero. Lastly, we note that some EoS predict more than one stable branch in the $M - R$ sequence and that we assign a prior probability of zero to all central densities corresponding to an unstable branch for a given EoS.

5. Perturbative QCD

Because of asymptotic freedom, QCD can be treated perturbatively at high densities ($\sim 40n_{\text{sat}}$) [152]. It has recently been argued that perturbative QCD offers nontrivial

constraints to the neutron star EoS when state-of-the-art next-to-next-to-next-to leading order (N3LO) perturbative results [14,153] are extended to neutron star densities using stability, causality, and consistency arguments [67,90]. The formalism was initially introduced in Ref. [90]. We briefly review its key components here but refer the reader to the original work for additional details.

Suppose an EoS can be characterized by a correlated set of values $\vec{\beta} = \{p(\mu), n(\mu), \mu\}$, where p is the pressure, n is the number density, and μ is the chemical potential. Also suppose that we have knowledge of the EoS at some low-density limiting value, μ_{low} , and a high-density limiting value, μ_{high} , meaning that we know

$$\vec{\beta}_{\text{low}} = \{p_{\text{low}}, n_{\text{low}}, \mu_{\text{low}}\} \equiv \{p(\mu_{\text{low}}), n(\mu_{\text{low}}), \mu_{\text{low}}\}, \quad (41)$$

$$\vec{\beta}_{\text{high}} = \{p_{\text{high}}, n_{\text{high}}, \mu_{\text{high}}\} \equiv \{p(\mu_{\text{high}}), n(\mu_{\text{high}}), \mu_{\text{high}}\}. \quad (42)$$

There are an infinite number of EoS that can connect $\vec{\beta}_{\text{low}}$ and $\vec{\beta}_{\text{high}}$, but any such EoS must respect thermodynamic stability, causality, and consistency. Thermodynamic stability implies that the grand-canonical potential from which the EoS is derived (Ω) is concave with respect to μ , meaning that $\partial_\mu^2 \Omega \leq 0$. At $T = 0$, we also have

$$\Omega(\mu) = -p(\mu), \quad (43)$$

$$n = \frac{\partial p}{\partial \mu}. \quad (44)$$

Therefore, stability results in a constraint on the slope of $n(\mu)$,

$$\frac{\partial n}{\partial \mu} \geq 0. \quad (45)$$

Moreover, the causality requirement constrains $c_s^2 \leq 1$, which, at $T = 0$, relates to $n(\mu)$ and $\partial_\mu n$ such that

$$c_s^{-2} = \frac{\mu \partial n}{n \partial \mu} \geq 1. \quad (46)$$

Combining stability and causality, we have that, at each point in $\mu - n$ space, the slope of the curve passing through that point corresponding to a maximally stiff ($c_s^2 = 1$) EoS is $\partial n / \partial \mu = n / \mu$.

Finally, because we must also ensure that at $(\mu_{\text{low}}, n_{\text{low}})$ the pressure is p_{low} and, similarly, that at $(\mu_{\text{high}}, n_{\text{high}})$ the pressure is p_{high} , it must also be true that

$$\int_{\mu_{\text{low}}}^{\mu_{\text{high}}} n(\mu) d\mu = p_{\text{high}} - p_{\text{low}} = \Delta p. \quad (47)$$

We can derive constraints on Δp based on stability and causality constraints on $n(\mu)$. We can place a lower bound on Δp by asking which curve connecting $(\mu_{\text{low}}, n_{\text{low}})$ to $(\mu_{\text{high}}, n_{\text{high}})$ minimizes the integral in Eq. (47) while still respecting stability and causality. We will call that quantity Δp_{min} . Equivalently, we can construct the curve which maximizes the integral in Eq. (47) and still respects stability and causality and denote that Δp_{max} . Assuming c_s^2 is bounded from above only by the causal limit, we have [90]

$$\Delta p_{\text{min}} = \frac{1}{2} \left(\frac{\mu_{\text{high}}^2}{\mu_{\text{low}}} - \mu_{\text{low}} \right) n_{\text{low}} \quad (48)$$

$$\Delta p_{\text{max}} = \frac{1}{2} \left(\mu_{\text{high}} - \frac{\mu_{\text{low}}^2}{\mu_{\text{high}}} \right) n_{\text{high}}. \quad (49)$$

These constraints combined imply that for any $\vec{\beta}_{\text{high}}$ and for a fixed $(\mu_{\text{low}}, n_{\text{low}})$, p_{low} must be between $[p_{\text{high}} - \Delta p_{\text{max}}, p_{\text{high}} - \Delta p_{\text{min}}]$.

These guidelines for connecting two arbitrary regimes via an EoS which respects stability, causality, and consistency can be used to extrapolate pQCD results to densities relevant to neutron stars. That is because, if we know $\vec{\beta}_{\text{high}} = \vec{\beta}_{\text{pQCD}}$, we can check if an EoS for which we only have knowledge up to a lower matching density $n_{\text{low}} = n_{\text{match}}$ can be connected to $\vec{\beta}_{\text{pQCD}}$ through a causal and stable EoS.

Our knowledge from pQCD is derived from current state-of-the-art calculations in Refs. [154, 155], which report a partial N3LO calculation of the zero-temperature,

high-density QCD grand-canonical potential. Because these results arise from a series expansion in the QCD coupling constant and are then truncated at a finite order, we have to estimate the error introduced by the missing higher-order (MHO) terms. In the case of QCD, the MHO error depends on a residual, unphysical renormalization scale, $\bar{\Lambda} \propto \mu$, which is underdetermined. Instead, the standard approach is to vary $\bar{\Lambda}$ around a fiducial scale by some fixed factor. We follow Ref. [67], which adopted a scale-averaging approach. That means that pQCD results are given as a family of independent predictions $\vec{\beta}_{\text{pQCD}}(X)$, where $X \equiv 3\bar{\Lambda}/2\mu_{\text{high}}$. We set $\mu_{\text{high}} = 2.6$ GeV based on Ref. [156], which points out that the uncertainty estimation for pQCD calculations at this value is similar to that of χ EFT at $1.1n_{\text{sat}}$ (about $\pm 24\%$ variation around the mean value [67]). We consider $X \in [1/2, 2]$, the same range that was implemented in Ref. [67] and that has been suggested by phenomenological models [157–160] as well as the large-flavor limit of QCD [161].

Now that we have defined the theoretical input from high densities, we need to discuss how we define the low-density input from GP and mGP EoS. For any neutron star EoS that we generate with the GP or the mGP, $\vec{\phi}_k$, we have to check that it can be connected to $\vec{\beta}_{\text{pQCD}}(X)$, for a given X , from $\vec{\beta}_{\text{low}} = \{p_k(n_{\text{match}}), n_{\text{match}}, \mu_k(n_{\text{match}})\}$. In practice, we check that $p_k(n_{\text{match}})$ leads to $\Delta p \in [\Delta p_{\text{min}}, \Delta p_{\text{max}}]$, given p_{high} from $\vec{\beta}_{\text{pQCD}}(X)$. Since the relevant scale for the neutron star EoS is the central density of a maximally massive star, n_B^{max} , we set $n_{\text{match}} = n_{B,k}^{\text{max}}$, which varies for each EoS. For the renormalization scale parameter, we follow Ref. [67] and sample 1000 values of $X \in [1/2, 2]$, evenly spaced in $\log(X)$. Hence, the pQCD weight associated with EoS k is

$$w_{\text{pQCD}}(\vec{\phi}_k) = \frac{1}{1000} \sum_{i=1}^{1000} \mathbf{1}_X(\vec{\phi}_k), \quad (50)$$

where $\mathbf{1}_X(\vec{\phi}_k)$ is the indicator function

$$\mathbf{1}_X(\vec{\phi}_k) = \begin{cases} 1, & \text{if } \Delta p_k \in [\Delta p_{\text{min}}, \Delta p_{\text{max}}] \\ 0, & \text{otherwise} \end{cases}, \quad (51)$$

and $\Delta p_k = p_{\text{pQCD}}(X) - p_k(n_{B,k}^{\text{max}})$. Recall that Δp_{min} and Δp_{max} can be obtained from Eqs. (48) and (49), using

$$\begin{aligned} \mu_{\text{high}} &= 2.6 \text{ GeV}, \\ \mu_{\text{low}} &= \mu_k(n_{B,k}^{\text{max}}), \\ n_{\text{high}} &= n_{\text{pQCD}}(\mu = 2.6 \text{ GeV}, X), \\ n_{\text{low}} &= n_{B,k}^{\text{max}}. \end{aligned}$$

Effectively, $w_{\text{pQCD}}(\vec{\phi}_k)$ captures how often, out of the 1000 values for X , $\vec{\phi}_k$ can be connected to $\vec{\beta}_{\text{pQCD}}(X)$ with an EoS that respects thermodynamic stability, causality, and consistency. This procedure defines a weighting scheme associated with input from pQCD, which suppresses EoS that are in tension with pQCD results by a factor proportional to the strength of the disagreement under the scale-averaging assumption.

We emphasize that the choice of n_{low} under this weighting scheme defines a termination density for the benchmark and modified GP priors on the EoS. Above n_{low} , which we take to be $n_{B,k}^{\text{max}}$, the EoS is only bounded by stability and causality. There are no assumptions about the smoothness of the EoS or the number of phase transitions allowed.⁷ Once a choice for n_{low} has been made, the posterior for the EoS is conditional on $n_B \leq n_{\text{low}}$ and should only be shown up to $n_B \leq n_{\text{low}}$. That is, although any benchmark or modified GP sample is defined up to a maximal central pressure of $\log_{10} p = 37 \text{ erg cm}^{-3}$, once pQCD input is applied at $n_{B,k}^{\text{max}}$, we impose that the EoS is only bounded by stability and causality at higher densities, and it would be inconsistent to show the GP EoS above $n_{B,k}^{\text{max}}$. For the interested reader, a detailed analysis on the interplay between imposing a prior above n_B^{max} and changing n_{low} was done in Ref. [162]. In this work, our choice of setting $n_{\text{low}} = n_{B,k}^{\text{max}}$ is motivated by the lack of observational/experimental constraints above the maximum central density realized in neutron stars. Although one could study a prior extending beyond this regime using pQCD input as a constraint (as was done in, e.g., Refs. [16, 67, 162, 163]), the posterior above n_B^{max} cannot not be verified by current or future observations of neutron stars. A promising direction beyond mass, radius, and tidal deformability is to study the gravitational waveform from binary neutron star mergers [164, 165] which could constrain the EoS at high densities and the maximum mass, and thus the maximum central density, of neutron stars.

6. Observational measurements not included in our analysis

We make the choice here to not include in our analysis other recent claims of very heavy or very light neutron stars. For example, there have been recent claims of pulsars heavier than the ones considered here, namely, PSR J1810 + 1744 at $2.13 \pm 0.04 M_{\odot}$ [166] and PSR J0952-0607 at $2.35 \pm 0.17 M_{\odot}$ [167], but possible systematic errors for these measurements are not as well understood as those involved in Shapiro time-delay-based measurements, such as those for PSR J0740 + 6620 and PSR J1614 – 2230. Specifically, the fit residuals in Ref. [166]

⁷Note that below $n_{B,k}^{\text{max}}$ the prior is determined by the choice of hyperparameters of the benchmark and modified GPs, which do constrain the smoothness and number of phase transitions in the EoS.

for the properties of the companion to spider-pulsar PSR J1810 + 1744 are clearly not a random scatter (see Fig. 1 in Ref. [166]), which suggests that the fit values and inferred mass are subject to systematic errors we do not currently understand. The picture is more promising for the inferred mass of PSR J0952-0607, where at least the residuals do not seem to indicate problems with the fit (see Fig. 1 in Ref. [167]). But there is still the question about whether the good fit indicates that the system is well understood from a theoretical perspective and whether the inferred mass is not just precise but also accurate.

In a separate measurement, the central compact object of the supernova remnant HESS J1731-347 was recently estimated to have a mass of $0.77_{-0.17}^{+0.20} M_{\odot}$ and radius of $10.4_{-0.78}^{+0.86} \text{ km}$ [168], possibly making it the lightest neutron star ever observed. Here, the low estimated mass stems from the use of a low distance to the source combined with the assumption that the surface radiates uniformly, which tends to favor a carbon atmosphere over a hydrogen or helium atmosphere. Moreover, in the fitting, it was assumed that surface magnetic fields can be ignored. However, nonuniform emission is consistent with data on several similar sources [169], making hydrogen and helium atmospheres possible and making it plausible that the neutron star in HESS J1731-347 could have a standard mass, well above $1 M_{\odot}$ [119]. Thus, to remain conservative on the data we use in this work, we do not consider PSR J1810 + 1744, PSR J0952-0607, and the center compact object in HESS J1731-347 in our analysis.

C. Model evidence

We have two distinct set of prior beliefs, or models, for the EoS. We combine samples from these two models into one prior, which we introduced in Sec. III C as Φ and which can be represented as the union of samples from the mGP and the benchmark GP, $\Phi = \Phi_{\text{mGP}} \cup \Phi_{\text{benchmark GP}}$.

We defined the evidence (\mathcal{E}) in Sec. IV A, where the integral in Eq. (18) is over all possible samples that can be generated from a model. In practice, we only have access to a finite number of samples, and \mathcal{E} is approximated as

$$\mathcal{E}_m \approx \sum_i^{N_m} \mathcal{L}(\vec{\phi}_k) q(\vec{\phi}_k), \quad \vec{\phi}_k \in \Phi_m, \quad (52)$$

where N_m is the number of samples from model $m = \{\text{benchmark GP, mGP}\}$, including the samples in $\Phi_m \cap \Phi_x$ for which $\mathcal{L}(\vec{\phi}_k) = 0$. Therefore, a key assumption is that we sample *enough* EoS to correctly approximate \mathcal{E} . This aspect further emphasizes the importance of the prior sample size checks we introduced in Sec. III C—if we do not have enough samples in the regions where $\mathcal{L}(\vec{\phi}_k)$ is non-negligible, we cannot correctly approximate \mathcal{E} .

Another assumption we make is that each EoS has an equal prior probability. This implies that Eq. (52) is now

$$\mathcal{E}_m \approx \frac{1}{N_m} \sum_i^{N_m} \mathcal{L}(\vec{\phi}_k), \quad \vec{\phi}_k \in \Phi_m. \quad (53)$$

That is, the evidence becomes a simple average over the likelihoods of all the EoS from model $m = \{\text{benchmark GP, mGP}\}$.

A reasonable objection to this assumption is to question whether the different hyperparameters in each framework should have been included explicitly as hyperpriors. In general, we expect that with an increased number of parameters we also increase our chances of describing the data but also that simplicity should be rewarded over complexity. Here, the mGP is more complex than the benchmark GP, so why is there not a penalty in the calculation of the model evidence for mGP EoS? Actually, the penalty exists, and it is included implicitly. To understand this argument, recall the discussion in Sec. III regarding assumptions that are implied in a non-parametric framework. We stated that our effective parameters are the value of the speed of sound at each sampled value of pressure, $c_s^2(p_i)$, and that the method and the hyperparameters we choose for generating $c_s^2(p_i)$ dictate both the prior distribution and the correlations across the *effective* parameter space. That means that the mGP covers a bigger space in terms of the possible functional forms for $c_s^2(p_i)$, which results in more bad predictions [$\mathcal{L}(\vec{\phi}_k) \approx 0$] compared to the simpler benchmark GP. Therefore, in order to be competitive with the benchmark GP and make up for a larger number of bad predictions, the mGP must be more accurate in describing the data than the benchmark GP. We do not need to include the different hyperpriors in Eq. (52) because (i) the benchmark GP and mGP hyperparameters are not the parameters being estimated, and (ii) the mGP model is implicitly penalized for its increased functional complexity in the effective parameter space because it covers a larger volume in that space where $\mathcal{L}(\vec{\phi}_k) \approx 0$ compared to the benchmark GP.

V. BAYESIAN ANALYSIS OF NONTRIVIAL FEATURES IN THE SPEED OF SOUND INSIDE NEUTRON STARS

Now that we have established methods for generating EoS that display long (benchmark GP) and multiscale (mGP) correlations in $c_s^2(n_B)$, we can implement the constraints discussed in Sec. IV and begin to answer specific physics questions from a Bayesian perspective.

We will begin with an important sanity check—does our new framework provide reasonable agreement with data, even when multiscale correlations and nontrivial features appear in $c_s^2(n_B)$? We will answer that question by looking at the mass-radius posteriors. Following a discussion of the mass-radius posteriors, we can explore other questions such as follows. Are the EoS posteriors sensitive to the structure

in $c_s^2(n_B)$? What is the maximum central density of a neutron star? Do new pQCD constraints have a strong influence on our analysis? Is a peak in $c_s^2(n_B)$ supported by existing constraints? Finally, is the GP or the mGP framework favored by the data?

To better understand our results, we use a plotting method for our priors and posteriors that, to our knowledge, has not been used to infer properties of the EoS in the literature. Let us first describe our approach for plotting the prior because there are subtle differences between its plotting method and that of the posterior. We bin our two-dimensional (variables X and Y , e.g., mass and radius or c_s^2 and n_B/n_{sat}) prior in bins of a certain width ΔX , ΔY . We denote a particular bin as a pair (X_i, Y_i) , such that a point (x, y) is in bin i if $X_i \leq x < X_i + \Delta X$ and $Y_i \leq y < Y_i + \Delta Y$. A given EoS k is characterized on the X-Y plane by a set of l total pairs of points $\{(x_k^1, y_k^1), \dots, (x_k^l, y_k^l)\}$, which produce a curve that passes through N_k of these bins. Every time the EoS passes through a bin i , we count that as a hit: $h_k(X_i, Y_i) = 1$. Otherwise, if the EoS does not pass through that location, we assign $h_k(X_i, Y_i) = 0$. We sum over all hits within a bin (i.e., count all the EoS that pass through it) to obtain the total hits for that specific bin, $\mathcal{H}_{\text{prior}}(X_i, Y_i)$, where

$$\mathcal{H}_{\text{prior}}(X_i, Y_i) = \sum_k^{N_{\text{EoS}}} h_k(X_i, Y_i), \quad (54)$$

and

$$h_k(X_i, Y_i) = \begin{cases} 1, & \text{if } X_i \leq x_k^{l^*} < X_i + \Delta X \wedge Y_i \leq y_k^{l^*} < Y_i + \Delta Y, \\ & l^* \in \{0, \dots, l\} \\ 0, & \text{otherwise.} \end{cases} \quad (55)$$

To normalize the hits within a bin, we determine the total number of hits across all of our EoS samples across all bins,

$$\mathcal{H}_{\text{tot}} = \sum_i \mathcal{H}_{\text{prior}}(X_i, Y_i), \quad (56)$$

such that our normalized distribution for the prior within a given bin (X_i, Y_i) is

$$\mathcal{N}_{\text{prior}}(X_i, Y_i) = \frac{\mathcal{H}_{\text{prior}}(X_i, Y_i)}{\mathcal{H}_{\text{tot}}}. \quad (57)$$

Our plotting method is similar for the posterior. However, because the posterior probability is proportional to the prior probability multiplied by the likelihood, we have to include the likelihood, $\mathcal{L}(\vec{\phi}_k)$, of each EoS when calculating the hits of the bin. More concretely,

$$\mathcal{H}_{\text{posterior}}(X_i, Y_i) = \sum_k^{N_{\text{EoS}}} \mathcal{L}(\vec{\phi}_k) h_k(X_i, Y_i), \quad (58)$$

and $\mathcal{N}_{\text{posterior}}$ has the same form as Eq. (57),

$$\mathcal{N}_{\text{posterior}}(X_i, Y_i) = \frac{\mathcal{H}_{\text{posterior}}(X_i, Y_i)}{\mathcal{H}_{\text{tot}}}, \quad (59)$$

but we use the posterior values instead of prior values to calculate \mathcal{H}_{tot} .

The procedure described above for representing prior and posterior distributions is equivalent to showing the prior as a normalized two-dimensional (2D) histogram, and the posterior as a weighted, normalized 2D histogram. We will refer to these plotting methods as the binned joint prior/posterior probability density. Representing probability densities in this way is advantageous as long as the bin sizes are chosen appropriately because we do not have to rely on kernel density estimates that can smear out important information. Kernel density estimates also perform poorly near sharp boundaries and may predict a finite probability density in regimes where data do not exist or, in cases where such constraints exist, beyond physical boundaries (i.e., $c_s^2 < 0$ or $c_s^2 > 1$). The only aspect that requires care when applying this method is the interplay between the size of the bins, $(\Delta X, \Delta Y)$, and the number of samples. Using a fine grid with too few samples can result in bin heights that fluctuate significantly between neighboring points. Similarly, using a grid that is too coarse when there are an adequate number of samples risks smearing out important features in the distribution, as can be the case with kernel density estimates.

Finally, the prior and posterior distributions we show in the sections below are conditional on $n_B \leq n_B^{\text{max}}$. That is, we only plot results for the EoS of an isolated, stable, and slowly rotating neutron star. We make this choice because the constraints used in this work only inform the posteriors up to n_B^{max} . Beyond that density, which is different for each EoS, the posterior will be dominated by the prior. The implications of this choice for interpreting the priors and posteriors shown here are discussed in the sections below.

A. Are mass-radius posteriors sensitive to structure in $c_s^2(n_B)$?

To fully explore the phase space of the EoS of neutron stars, it is important to have a broad prior in the mass-radius relation. In the left panels of Fig. 4, we show the stable branches of the mass-radius prior for the benchmark GP (top) and the mGP (bottom) using the plotting method described in Eqs. (54)–(56). Here, we only plot the priors corresponding to samples in the set that meets the basic checks we outlined in Sec. III C (Φ_{v}). We find that both priors produce a similarly diverse set of mass-radius curves, with the highest prior regimes passing through $R = 10\text{--}14$ km and up to masses $M \sim 2\text{--}2.2M_{\odot}$. We also

see that both priors allow for maximum masses up to $M \sim 3.5M_{\odot}$, although the prior disfavors $M \gtrsim 2.2M_{\odot}$. White regions in Fig. 4 indicate that no samples in Φ_{v} reach that region in the $M - R$ plane. The bottom right-hand side of the left panels in Fig. 4—the large-radius, low-mass regime—is primarily ruled out by constraints on the symmetry energy, which predict a soft EoS in that regime. The top left-hand side of the left panels in Fig. 4—the small-radius, high-mass regime—is ruled out mostly because it is beyond the point of stability for the $M - R$ sequences that reach such high masses. Comparing our priors for the benchmark GP and the mGP, we find they are nearly identical, despite the significant differences in how the EoS are constructed.

After applying the constraints outlined in Sec. IV, we then produce our posteriors, which are shown on the right panels of Fig. 4. We use the plotting technique described in Eqs. (54)–(56) to present the posterior distribution for the benchmark GP (top panels) and the mGP (bottom panel) models. In addition, in Table II, we present the median, 68%, and 90% credible regions from the benchmark GP and mGP posteriors for the equatorial radius of a $1.4M_{\odot}$ star ($R_{1.4}$), the equatorial radius of a $2.1M_{\odot}$ star ($R_{2.1}$), and maximum stellar mass (M_{max}), assuming an isolated, slowly rotating star. We find that in terms of the observables highlighted in Table II as well as the binned joint posteriors the results for the benchmark GP and the mGP are nearly identical.

Observe that both benchmark GP and mGP EoS support neutron star masses up to $M \sim 2.7M_{\odot}$ (albeit with a smaller likelihood), but neutron stars with higher masses are extremely disfavored. Observe also that, from the credible bands alone, one would reasonably assume that all neutron stars heavier than $M \sim 2.7M_{\odot}$ must have larger radii at high masses. However, from the joint posterior, it is clear that a number of mass-radius curves for heavy neutron stars may either be nearly straight or even bend slightly to the left. In any case, we see that, although specific M–R curves may be affected by structure in the speed of sound, the latter does not affect the M–R posteriors, given currently available observations.

B. Are EoS posteriors sensitive to structure in $c_s^2(n_B)$?

Next, we test if the different prior assumptions made about correlations across densities in the benchmark GP versus the mGP model lead to any significant differences in the posterior for $c_s^2(n_B)$. Recall that the benchmark GP produces smooth c_s^2 curves with uniform correlations across densities, whereas the mGP display sharp features in c_s^2 and multiscale correlations across densities (refer also to Figs. 1 and 3).

In Fig. 5, we show the priors (left panels) and the posteriors (right panels) for c_s^2 as a function of n_B in units of n_{sat} , up to n_B^{max} for each EoS. Again, benchmark GP and mGP results are shown on the top and bottom panels,

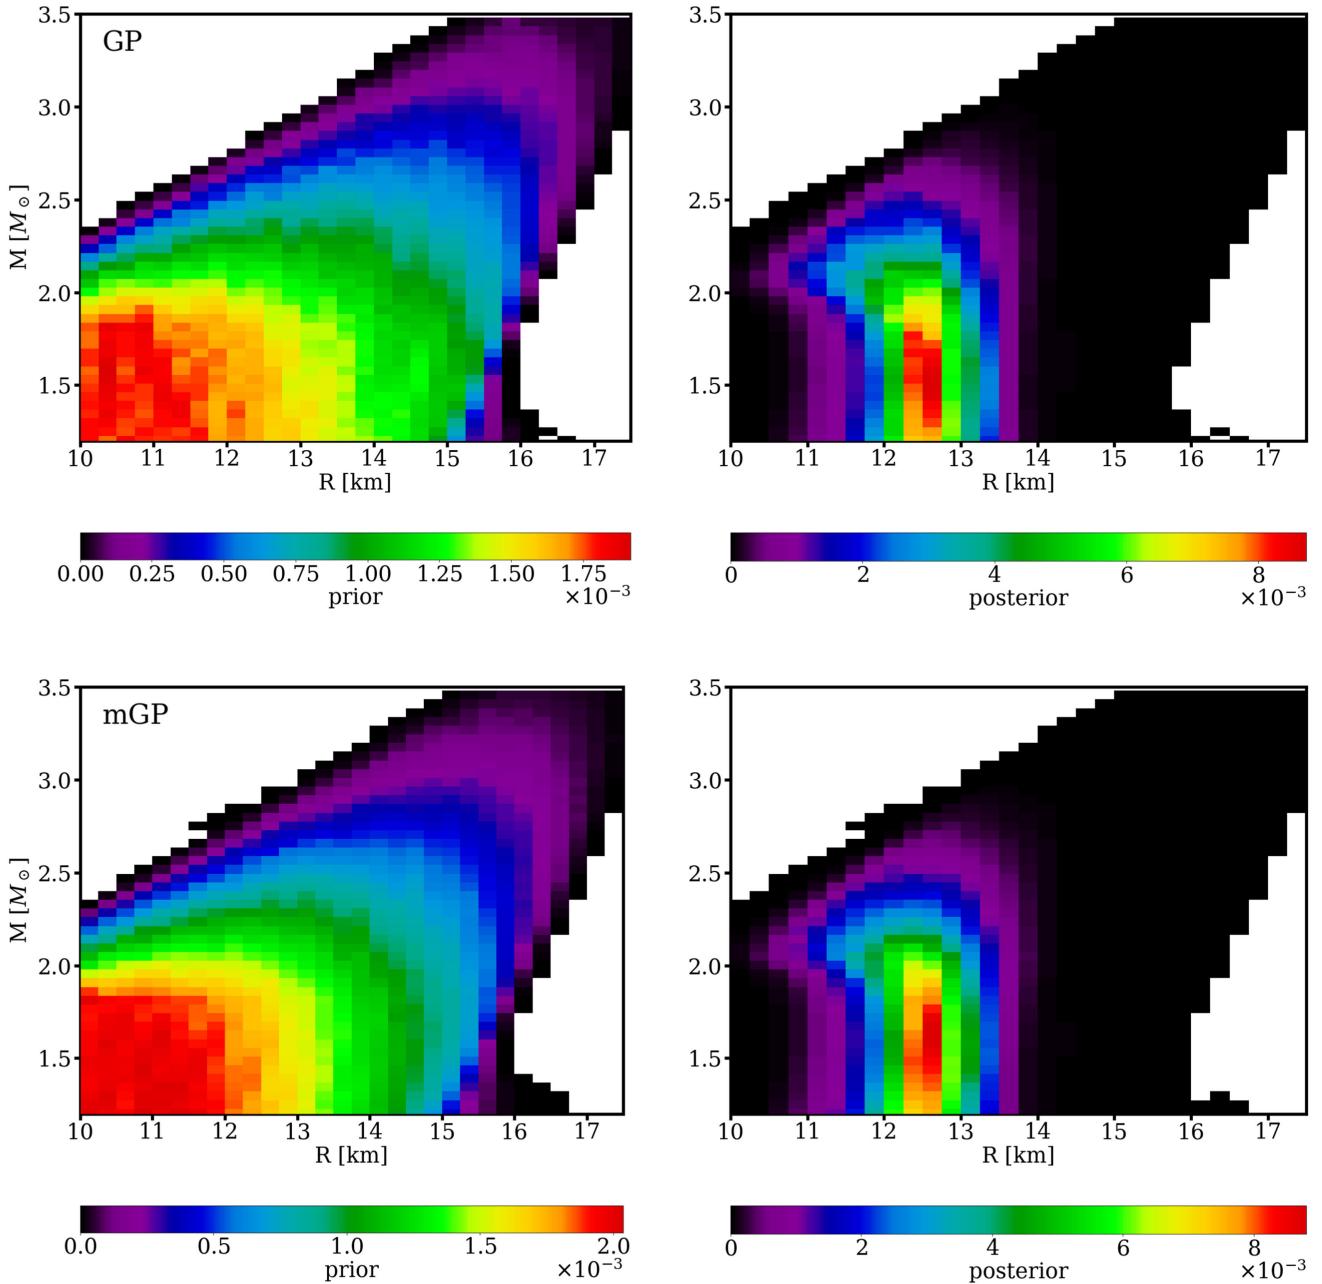


FIG. 4. Mass-radius prior (left) and posterior (right) probability distributions for GP (top) and mGP (bottom) EoS. Both the prior and posterior probability distributions are produced by binning the EoS by mass and radius and then normalizing the heights of the bins such that their sum is equal to 1. For the posteriors, each EoS is weighted by the corresponding likelihood. Observe that the joint posteriors are similar when using the GP or mGP models, with masses larger than $2.7M_{\odot}$ and radii larger than 14 km disfavored.

TABLE II. Median, 68%, and 90% credible regions from the benchmark GP and mGP posteriors for the equatorial radius of a $1.4M_{\odot}$ star ($R_{1.4}$), the equatorial radius of a $2.1M_{\odot}$ star ($R_{2.1}$), and maximum stellar mass (M_{\max}), assuming an isolated, slowly rotating star.

EoS	$R_{1.4}$ (km)	$R_{2.1}$ (km)	M_{\max} (M_{\odot})
Benchmark GP	12.55, (12.07, 13.05), (11.56, 13.37)	12.26, (11.52, 12.90), (11.02, 13.34)	2.26, (2.10, 2.51), (2.04, 2.68)
mGP	12.55, (12.04, 13.04), (11.54, 13.33)	12.34, (11.62, 12.98), (11.12, 13.39)	2.25, (2.09, 2.50), (2.03, 2.67)

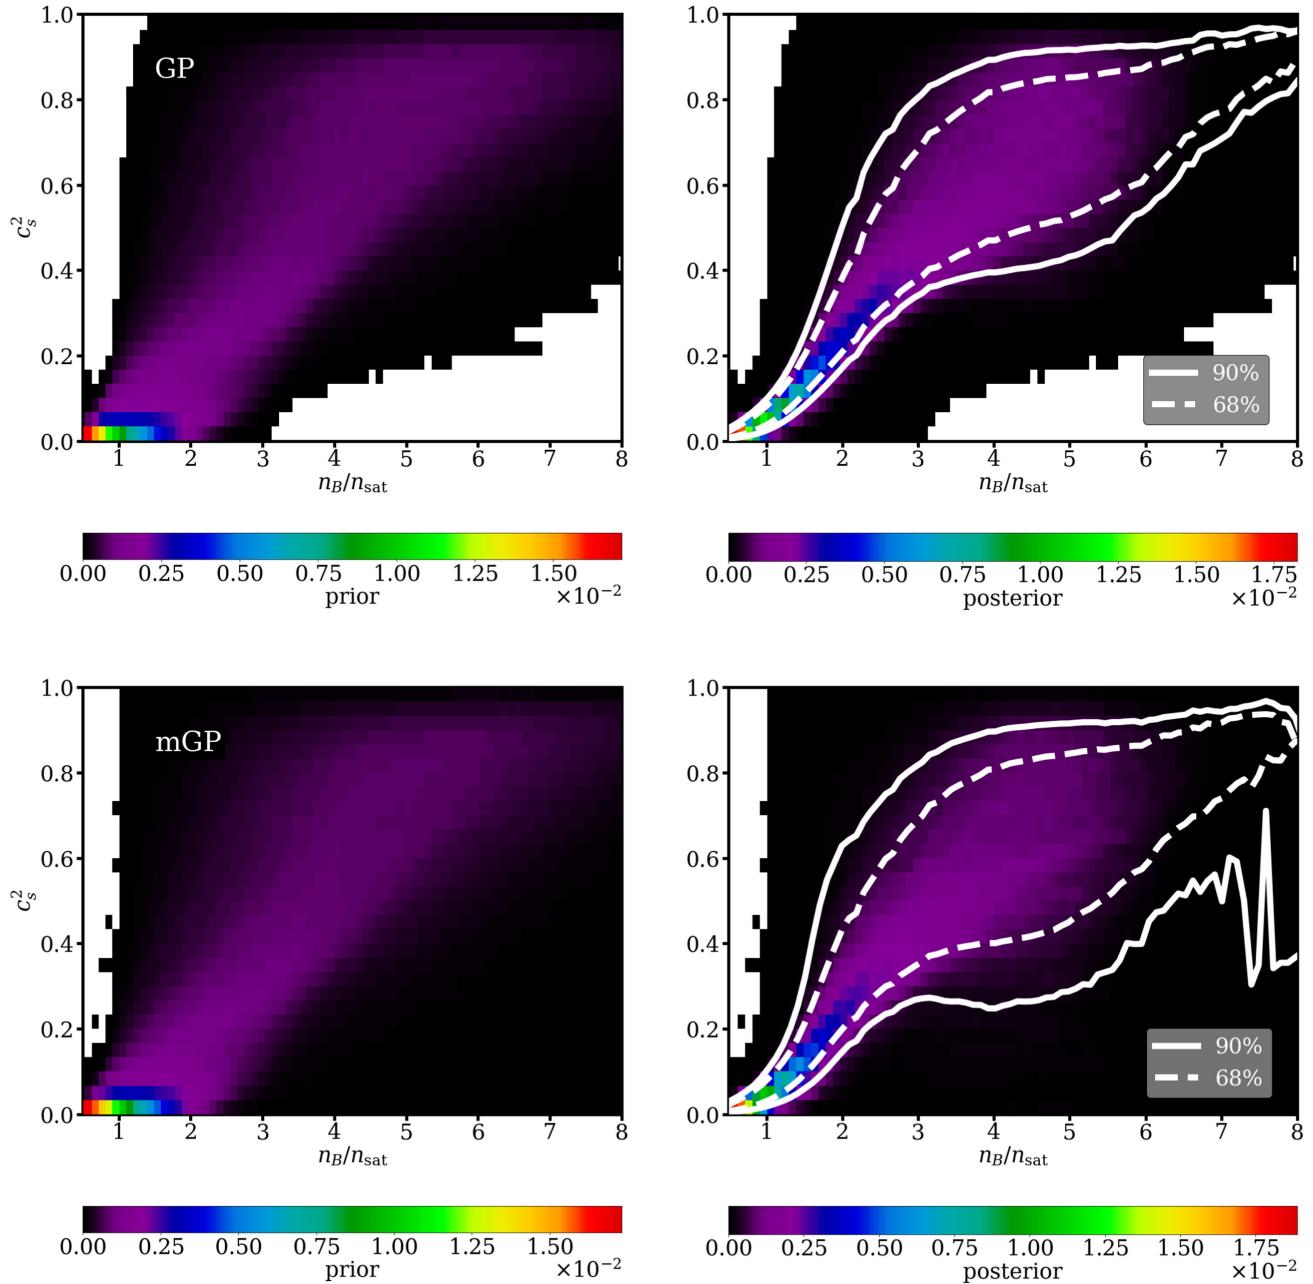


FIG. 5. EoS prior (left) and posterior (right) probability distributions for GP (top) and mGP (bottom) EoS conditional on $n_B \leq n_B^{\max}$ for each EoS. The EoS are represented by the speed of sound squared in units of c^2 as a function of baryon number density in units of n_{sat} . The prior and posterior probability distributions are produced by binning the EoS by the speed of sound and number density and then normalizing the heights of the bins such that their sum is equal to 1. For the posteriors, each EoS is weighed by the corresponding likelihood. Also shown in the posterior plots are the 90% and 68% credible regions for the speed of sound squared at a given density for $0.5 \leq n_B \leq 8.0n_{\text{sat}}$. The posterior probability that the central density for a maximally massive star is greater than $\sim 6.0n_{\text{sat}}$ is negligible in both cases. We note that probability densities are low in the regime between 2 and $6n_{\text{sat}}$ because of the wide spread in the allowed behavior for c_s^2 . Observe that at the 90% level the mGP posterior is wider than the GP one for all densities above n_{sat} .

respectively. The priors are shown using the binning method outlined in Eqs. (54)–(56). The posteriors are shown both in terms of the binned joint posteriors as well as constant density 68% and 90% credible bands. In these plots, we only show c_s^2 up to the maximum central baryon density for a stable star in the slow-rotation regime, n_B^{\max} .

Most EoS lead to n_B^{\max} around $4 - 7n_{\text{sat}}$ with a handful that extend up to $8n_{\text{sat}}$. Thus, we plot only up to $n_B = 8.0n_{\text{sat}}$. We will discuss the posterior for n_B^{\max} separately in Fig. 6. We note that Fig. 5 only includes EoS in Φ_{\vee} (the set of EoS that pass the constraints discussed in Sec. III), which clearly affects the priors. This selection leads to $c_s^2(n_B)$ functional

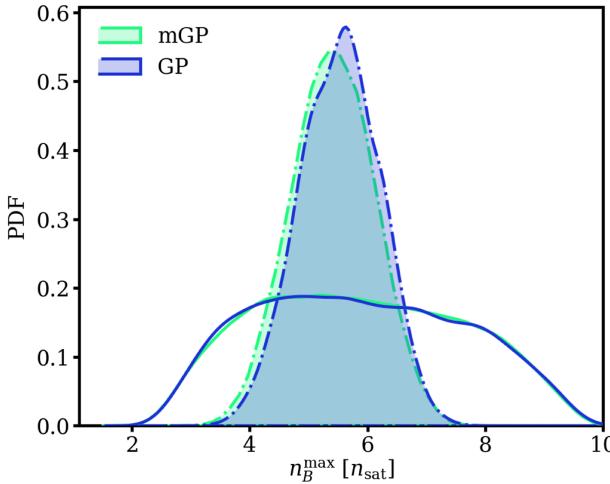


FIG. 6. The estimated prior (solid lines) and estimated posterior (dot-dashed, filled lines) probability density distributions for the maximal central density of a stable, nonrotating star in units of $n_{\text{sat}}^{\text{max}}$ for EoS from the benchmark GP (green) and mGP (blue). The priors for GP and mGP produce nearly identical PDFs for n_B^{max} such that the lines are indistinguishable from each other. The vast majority of EoS from the posterior predict a maximal central density between ~ 4 – $8n_{\text{sat}}$.

forms that favor high values ($c_s^2 \geq 1/3$) at large densities ($n_B \geq 2n_{\text{sat}}$).

Let us first discuss the priors. Observe that the highest probability regions look nearly identical between the benchmark GP and mGP models when we only consider EoS in Φ_{\checkmark} . In both cases, at very low densities ($n_B \leq n_{\text{sat}}$), there is a strong preference for a nearly vanishing c_s^2 . This result is unsurprising because we use the same crust at low densities for both models and also impose that no modifications are introduced in c_s^2 below $n_B < 1.1n_{\text{sat}}$ for mGP EoS. Progressing to intermediate densities ($n_{\text{sat}} \leq n_B \leq 3n_{\text{sat}}$), we find a general trend in both priors to larger c_s^2 , but this trend has a rather wide spread, allowing for diverse behavior in c_s^2 , hence the lower probability within the credible bands. At densities above $3n_{\text{sat}}$, we find a general trend upward in c_s^2 , but again with an even larger spread.

One key difference, however, does exist between the benchmark GP and mGP priors. The benchmark GP EoS in Φ_{\checkmark} do not contain *any* samples that have a low c_s^2 at large n_B (notice the large white region in the bottom-left corner of the top-left panel in Fig. 5). On the other hand, the same region in the mGP prior has a nonzero prior probability density. The key difference is that the benchmark GPs are smooth and domain points are correlated over a long range in density. Therefore, the speed of sound functional forms from the benchmark GP cannot easily fluctuate downward to this region (especially since they need to support neutron stars with $M \geq 1.8M_{\odot}$). In contrast, functional forms from the mGP model can have fluctuations to larger c_s^2 , followed by regions of lower c_s^2 . In this way, it is clear that the mGP

model allows us to explore a wider regime in EoS parameter space.

Next, let us discuss the binned joint $n_B - c_s^2$ posteriors and the constant density 90% and 68% credible regions, shown in the right panels of Fig. 5. The posteriors and credible bands are similar between the benchmark GP and the mGP, but they are not identical. We find that, at the 90% level, the mGP posterior is wider than the GP one for all densities above n_{sat} . Notably, the mGP posterior allows for slightly stiffer EoS in the regime $1.5 \lesssim n_B \lesssim 3.0n_{\text{sat}}$ and slightly softer EoS above $3.0n_{\text{sat}}$. For instance, at twice nuclear saturation density and 68% credibility, we extract $c_s^2 = 0.29_{-0.11}^{+0.27}$ using EoS from the GP posterior and $c_s^2 = 0.29_{-0.14}^{+0.34}$ using the mGP posterior. At four times nuclear saturation density, the GP EoS range is $c_s^2 = 0.63_{-0.23}^{+0.27}$, while the mGP EoS allows for $c_s^2 = 0.59_{-0.34}^{+0.31}$. Thus, we find that the GP leads to slightly stiffer EoS and has slightly narrower posterior credible bands. As we will discuss later in Fig. 6, very few EoS reach beyond $n_B > 6n_{\text{sat}}$, so the statistics in that regime are not sufficient to draw conclusions about differences between the two frameworks. This is evident from the highly oscillatory behavior of the mGP credible bands in that region.

We can draw further conclusions from the binned joint probability density posteriors. In both posteriors, there is a strong preference (blue regions) for a sharp rise in the speed of sound between 1 and $2n_{\text{sat}}$. Furthermore, at densities larger than $2n_{\text{sat}}$, large c_s^2 is favored, well beyond the conformal limit of $c_s^2 = 1/3$. However, $c_s^2(n_B)$ is significantly less constrained in that regime. Lastly, we also see from the binned joint posteriors that a large fraction of EoS must reach their maximum central densities around $\sim 5n_{\text{sat}}$ because the probability density decreases significantly beyond that point.

While the relative differences in the posteriors are not huge between the GP and mGP models, the point still holds that when the c_s^2 is allowed to display sharp features the posteriors are wider than when a smooth EoS is presumed. This is because implicitly imposing smoothness in the EoS through benchmark EoS translates to a prior that disallows low speeds of sound at high densities. We, therefore, argue that such sharp features should be adequately represented in priors for the extraction of the EoS in neutron star regimes. Furthermore, it is clear that if one is especially interested in studying whether low values of c_s^2 at high densities are allowed by nature it is even more important to allow for sharp features in c_s^2 .

C. Are maximum central density posteriors sensitive to structure in $c_s^2(n_B)$?

We will now discuss the priors and posteriors for the maximal central baryon density reached by a stable neutron star in the slow rotation regime, n_B^{max} . From Fig. 5, we already saw hints that n_B^{max} must peak around $5n_{\text{sat}}$ because

the posterior probability densities are low at baryon densities higher than that. To study this systematically, we have plotted the range of n_B^{\max} for both our priors and posteriors in Fig. 6 using a kernel density estimate.⁸

Observe that the n_B^{\max} priors for the benchmark GP (blue line) and mGP (green line) are essentially identical and mostly overlapping, ranging between $n_B^{\max} = 2-10n_{\text{sat}}$ (again, recall that our priors shown here represent only the samples in $\Phi_{\mathcal{V}}$). However, there is a significant change in the posterior probability density for both the benchmark GP and the mGP compared to their priors. Both posteriors are sharply peaked at $n_B^{\max} = 5-6n_{\text{sat}}$ with a range between $n_B^{\max} = 4-8n_{\text{sat}}$ (with essentially no EoS that produce $n_B^{\max} > 8n_{\text{sat}}$). The posterior for the mGP model peaks at a slightly smaller n_B^{\max} than that of the GP model, but the difference between the two distributions is very small. These results are consistent with what is shown in Fig. 5, in which the probabilities are compatible with zero for $n_B \geq 8n_{\text{sat}}$.

D. What is the impact of pQCD constraints?

As discussed in Sec. IV B 5, in order to include pQCD constraints, one must choose a matching density, n_{matching} . This matching density reflects how pQCD results are being propagated via stability, causality, and consistency constraints. In this work, we used $n_{\text{matching}} = n_{B,k}^{\max}$, the maximal central density predicted by EoS k for a stable, nonrotating neutron star. Although $n_{B,k}^{\max}$ changes for each EoS, we saw in Sec. V C and Fig. 6 that the posterior probability density for n_B^{\max} drops sharply above $\sim 6n_{\text{sat}}$ and is negligible beyond $8n_{\text{sat}}$.

In Fig. 7 (top), we quantify how pQCD constraints affected our inference of the EoS. In the bar chart, we show the proportion of EoS in our prior that are consistent ($w_{\text{pQCD}} = 1$ exactly), in tension ($0 < w_{\text{pQCD}} < 1$), and inconsistent ($w_{\text{pQCD}} = 0$ exactly) with pQCD input [see Eq. (50) for the definition of w_{pQCD}]. A vast majority of the EoS in the prior ($\sim 96.5\%$) are consistent with pQCD results, $\sim 3.5\%$ are in tension, and only 0.0083% (11 total) EoS are inconsistent with pQCD results. Because pQCD only completely excludes a very small fraction of EoS, these constraints cannot contribute strongly to the $M - R$ and $c_s^2(n_B)$ posteriors, shown in the previous section.

We observe, however, that a non-negligible population ($\sim 3.5\%$) that is in tension with pQCD results indeed exists. In Fig. 7 (bottom), we show a histogram for w_{pQCD} only for the samples that were found to be in tension with pQCD. Recall that w_{pQCD} is proportional to the strength of the disagreement between a given EoS and pQCD results over the range of $\bar{\Lambda}$ we sampled. Thus, we can quantify the

⁸For this observable, the probability densities vary smoothly with n_B^{\max} , and hence the use of kernel density estimates instead of the binning method we discussed earlier is safe.

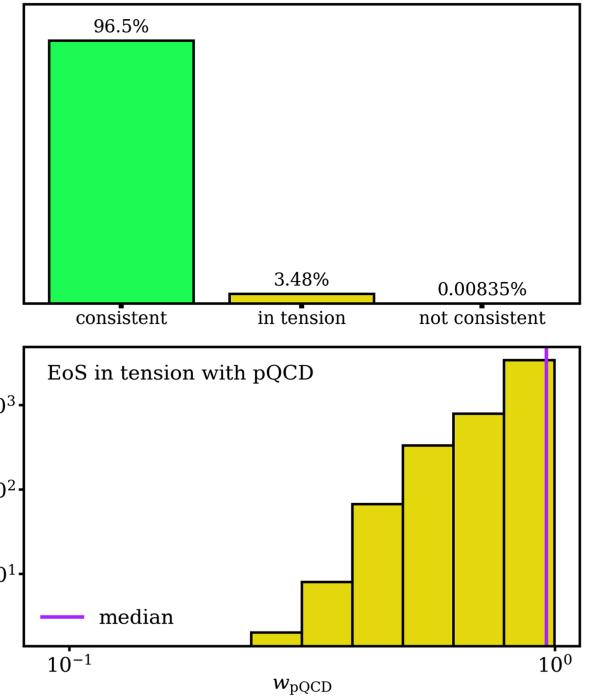


FIG. 7. Top: bar chart of the percentage of EoS that are consistent ($w_{\text{pQCD}} = 1$), in tension ($0 < w_{\text{pQCD}} < 1$), and not consistent ($w_{\text{pQCD}} = 0$) with pQCD input based on the formalism in Ref. [67]. Bottom: histogram of the w_{pQCD} for the 4592 EoS ($\sim 3.5\%$ of the total number of samples) in the combined benchmark GP and mGP prior that are in tension with pQCD input. The impact of pQCD input when $n_{B,k}^{\max}$ is used as the matching density is negligible (see Sec. IV B 5 for definitions).

impact of these samples on the posterior probability by looking at their w_{pQCD} distribution. As shown in Fig. 7, the median of the distribution is ~ 0.95 , meaning that more than half of the samples in tension with pQCD are only marginally suppressed.

The results shown here suggest that the impact of pQCD input on the inference of the neutron star EoS is minimal. Given that the opposite has been reported in the literature recently [16,67,170], we would like to address why that is the case. Let us recall Sec. IV B 5, in which we discussed the assumptions associated with including pQCD input in the analysis. We assume we know the EoS at a low-density limit [see Eqs. (41) and (42)], each determined by a set of three values that fix the EoS in those limits. These values correspond to the number density (n), the chemical potential (μ), and the pressure (p) at each limit. In the low-density regime, we get μ_{low} and p_{low} from the EoS, but we need to make a choice for n_{low} . This choice is important because $n_{\text{low}} = n_{\text{matching}}$, meaning that pQCD results will be propagated down to n_{low} . This choice is in principle arbitrary, but given that the largest density scale relevant to the EoS of isolated, slowly rotating neutron stars is n_B^{\max} , it is reasonable to impose $n_{\text{low}} = n_B^{\max}$ for each EoS. This is the choice that we make in this work and the choice that was

made in Ref. [171], which also reported that pQCD only affected a very small number of EoS in the prior.

Works that found a robust softening of the EoS at high densities, leading to a peak in the $c_s^2(n_B)$ posterior, used $n_{\text{matching}} = 10 n_{\text{sat}}$, where pQCD is more constraining because it is closer to $\mu_{\text{high}} = 2.6$ GeV. However, as shown in Fig. 6, the n_B^{max} posterior is essentially zero for $n_B^{\text{max}} > 7 n_{\text{sat}}$. That means that in Refs. [16,67,170] pQCD constraints are being imposed far beyond the densities where most realistic EoS predict a stable neutron star. There are no constraints from astronomical observations in those densities, so the impact of pQCD on the posterior will depend on prior-imposed assumptions about the correlations in $c_s^2(n_B)$ in the regime above n_B^{max} . We note here, as Refs. [172,173] also noted, that such results are very sensitive to the prior.

Imposing pQCD constraints at n_B^{max} with $X = [1/2, 2]$ results in only 3.5% of the EoS being affected. Only 11 are completely ruled out. Again, we highlight that in Fig. 7, which shows a histogram of the pQCD weights assigned to EoS that were suppressed by pQCD, we see that the vast majority of the EoS affected were only marginally suppressed. With these results, we conclude that our posteriors are dominated by astrophysical observations, which is why we do not see a softening of the EoS at larger densities. Nonetheless, pQCD offers nontrivial constraints even when incorporated exclusively at densities where an EoS predicts stable, slowly rotating neutron stars exist. Lastly, we point out that our findings are in agreement with those of Ref. [171], which found that pQCD affects the EoS mainly beyond the densities realized in neutrons stars.

E. Does $c_s^2(n_B)$ display a peak within neutron star densities?

Given that mGP EoS lead to reasonable mass-radius and $c_s^2(n_B)$ posteriors, we can now begin to explore the existence of structure in $c_s^2(n_B)$. One way to study the latter is to look for a bump [i.e., $c_s^2(n_B)$ rises and reaches a global maximum at some $n_B < n_B^{\text{max}}$ before decreasing again] that would signify a crossover phase transition. This type of structure is being actively discussed in the literature [16,18,19,67,170,174,175] as a signature of quark matter in the core of massive neutron stars. Such a structure has become especially relevant after studies that use pQCD constraints applied beyond densities realized in most realistic neutron star EoS found a posterior for $c_s^2(n_B)$ that displays a peak within neutron star densities [16,67].

One caveat here is that, as seen in Fig. 1, $c_s^2(n_B)$ can oscillate or contain first-order phase transitions. Thus, the *absence* of a global maximum in $c_s^2(n_B)$ before n_B^{max} does not imply that a phase transition does not occur within neutron star densities. Similarly, the *presence* of a global maximum in $c_s^2(n_B)$ before n_B^{max} does not confirm a transition to quark matter occurs in the core of massive

neutron stars because, as shown in Fig. 1, the onset of degrees of freedom other than quarks, such as heavy resonances or hyperons, can also cause the EoS to soften in a similar way. With this caveat in mind, we define the density at which the maximum in c_s^2 is reached as $n_B(c_{s,\text{max}}^2)$, or n_B at maximum c_s^2 .

The left panel of Fig. 8 shows the prior for $n_B(c_{s,\text{max}}^2)$ for both the benchmark GP and mGP models. The maximum value of c_s^2 can occur at any density up to $n_B \leq 10n_{\text{sat}}$, as seen for both priors. However, already at the level of the prior, we do see differences between the two models. The benchmark GP model prior has a peak at approximately the same density as the n_B^{max} posterior shown in Fig. 6. This feature indicates that benchmark GP samples mostly reach a global maximum in c_s^2 at or near their maximal baryon density, i.e., $n_B(c_{s,\text{max}}^2) \approx n_B^{\text{max}}$. This result is compatible with the benchmark GP assumption that the low-density (below $\sim 1.1 n_{\text{sat}}$) and high-density (above $\sim 1.1 n_{\text{sat}}$) regions display correlations of length $\ell = 1$ in units of $\log_{10} \text{erg cm}^{-3}$. The EoS in the low-density regime must be smooth and soft to be in agreement with symmetry energy estimates, while astronomical observations require an EoS stiff enough to support $2M_\odot$ stars but not too stiff in the regime below $3n_{\text{sat}}$ because of tidal deformability constraints around $1.4M_\odot$ stars. This transition from soft to stiff, by construction, happens over a range in pressure corresponding to the hyperparameter ℓ . Since we imposed a reasonably large value for ℓ , bumps are less likely in the benchmark GP model. In contrast, the mGP model has a prior that is relatively uniform in the range between $1.1 - 8n_{\text{sat}}$, rather than peaked at densities above $3n_{\text{sat}}$, as in the benchmark GP case. That implies that some of these EoS have a low-density bump in c_s^2 . In this case, the assumption that the low- and high-density regimes are correlated over a long range in pressure is relaxed, allowing for low-density bumps to appear.

In the right panel of Fig. 8, we show the posterior for the density at which the speed of sound is maximal (computed from a kernel density estimate), together with the prior, for both the benchmark GP (blue) and mGP models (green). The posteriors present interesting features and striking differences between both models. The benchmark GP model leads to a posterior with a maximum consistent with its prior and, thus, consistent with the maximum density n_B^{max} . However, the mGP posterior distribution for n_B at the global c_s^2 maximum is *bimodal*, with peaks at $\sim 2n_{\text{sat}}$ and $\sim 5n_{\text{sat}}$. The peak of the mGP posterior distribution centered at $n_B \sim 5n_{\text{sat}}$ is somewhat larger than the peak at $n_B \sim 2n_{\text{sat}}$. Nonetheless, the peak at $n_B \sim 2n_{\text{sat}}$ is still quite significant, and it is clearly a result of the extra structure built into the mGP model. We note that this low-density bump is consistent with recent preliminary findings from heavy-ion collisions [176] and may be an indication of a crossover phase transition (a possible explanation is quarkyonic matter; see, e.g., Ref. [24]). We should be careful in our

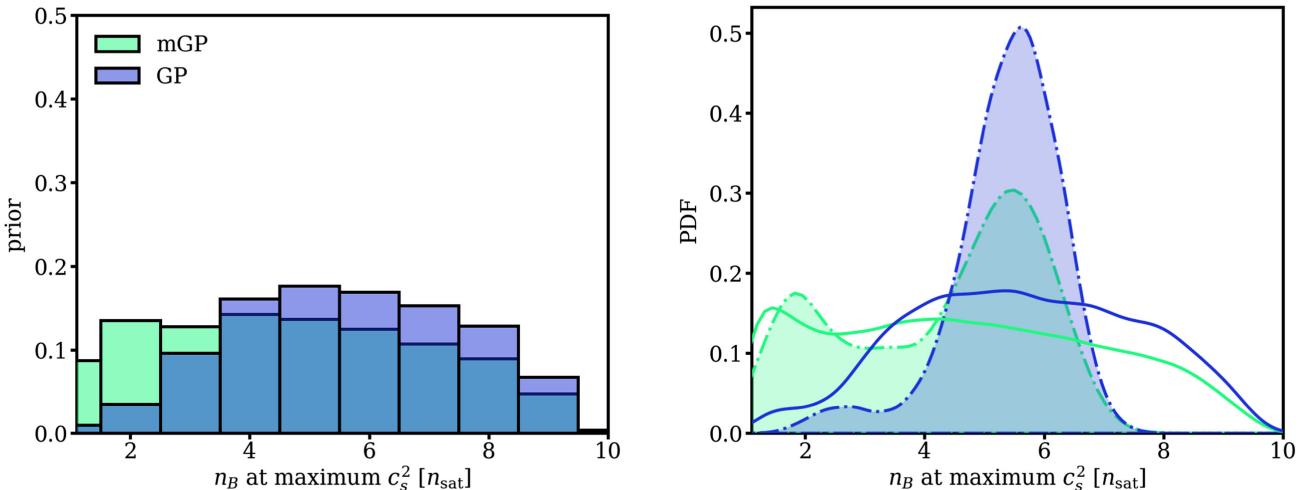


FIG. 8. Left: histograms of the prior distributions for the GP and mGP EoS of the value of the baryon density n_B in units of n_{sat} at the global speed of sound maximum for a stable nonrotating neutron star. The increased flexibility of mGP EoS allows for a global c_s^2 maximum to occur at lower densities compared to the GP EoS. Right: the estimated prior probability density distributions corresponding to the histograms on the left (solid curves) and estimated posterior probability density distribution after constraints on the symmetry energy, mass, radius, and binary tidal deformability are imposed along with pQCD input (dot-dashed curves). Data and theoretical constraints do not rule out a global c_s^2 maximum at densities below $3n_{\text{sat}}$ but are also not yet informative enough to favor or disfavor it over a global c_s^2 maximum above $3n_{\text{sat}}$.

interpretation of these results, however, since the posterior is bimodal, and a monotonically increasing c_s^2 cannot be ruled out. It remains to be seen if the low-density peak in the posterior for $n_B(c_s^2, \text{max})$ will be further enhanced or suppressed by future astrophysical observations of neutron stars.

We now investigate the differences between the two distinct peaks in the mGP posterior shown in Fig. 8 by dividing the mGP EoS into two groups: one with $n_B(c_s^2, \text{max}) < 3n_{\text{sat}}$ and one with $n_B(c_s^2, \text{max}) > 3n_{\text{sat}}$. In Fig. 9, we show the mGP posterior for $c_s^2(n_B)$ in the first group (left panel) and in the second group (right panel), using the binned probability densities from Eqs. (54)–(56), together with the constant density 90% and 68% credible regions. The resulting posteriors are qualitatively different between the two groups. We see a much sharper increase in c_s^2 at low densities for the $n_B(c_s^2, \text{max}) < 3n_{\text{sat}}$ group in the left panel. In that case, we see that c_s^2 may have a peak, followed by a decrease, or it may plateau, at larger n_B . Interestingly, the $n_B(c_s^2, \text{max}) < 3n_{\text{sat}}$ group allows for the possibility of a softening in the EoS at large densities that is not seen in the other group. The $n_B(c_s^2, \text{max}) > 3n_{\text{sat}}$ group on the right panel appears to have a monotonically increasing posterior that ends at a large c_s^2 at large n_B . This group more closely resembles nucleonic-only EoS. Another interesting difference is that, unlike the $n_B(c_s^2, \text{max}) < 3n_{\text{sat}}$ group, the $n_B(c_s^2, \text{max}) > 3n_{\text{sat}}$ group has a tight c_s^2 distribution at large n_B , which drives c_s^2 to large values.

We can also analyze quantitative differences between these posteriors. In the left panel, at $2n_{\text{sat}}$, the

$n_B(c_s^2, \text{max}) < 3n_{\text{sat}}$ group allows for c_s^2 as high as ~ 0.80 at 90% credibility and ~ 0.55 at 68% credibility. In contrast, the second group predicts much smaller c_s^2 for $n_B = 2n_{\text{sat}}$, ~ 0.45 and ~ 0.35 at 90% and 68% credibility, respectively. At densities above $3n_{\text{sat}}$, the $n_B(c_s^2, \text{max}) < 3n_{\text{sat}}$ group continues to allow for a wide range of c_s^2 , displaying c_s^2 values as low as ~ 0.2 at 90% credibility and as high as ~ 0.8 at 90% credibility at $5n_{\text{sat}}$. At 68% credibility, the lower and upper bands are at roughly $c_s^2 \approx 0.3$ and $c_s^2 \approx 0.7$, respectively, at $5n_{\text{sat}}$. In contrast, at $5n_{\text{sat}}$, the ranges for the $n_B(c_s^2, \text{max}) > 3n_{\text{sat}}$ group are approximately $[0.5, 0.8]$ at 68% credibility and $[0.3, 0.9]$ at 90% credibility. Overall, if a global maximum occurs below $3n_{\text{sat}}$, our results indicate that we can expect an EoS that is stiffer at low densities and softer at high densities. On the other hand, if a global maximum occurs above $3n_{\text{sat}}$, the c_s^2 posterior suggests that the EoS is stiffer and above the conformal value of $1/3$ for all $n_B > 3n_{\text{sat}}$.

Recalling an earlier discussion about what the absence of a clear peaklike structure in c_s^2 means, we emphasize that the EoS in both posteriors shown in Fig. 9 were generated using the mGP framework. Therefore, all these samples contain nontrivial features. Thus, it is possible that the $n_B(c_s^2, \text{max}) > 3n_{\text{sat}}$ group may have a small bump in c_s^2 at low densities but then the EoS continues to become stiffer, ending at an even larger c_s^2 near n_B^{max} .

F. Are there nontrivial features in $c_s^2(n_B)$?

We have established that different assumptions about the scale of correlations across densities in the speed of sound

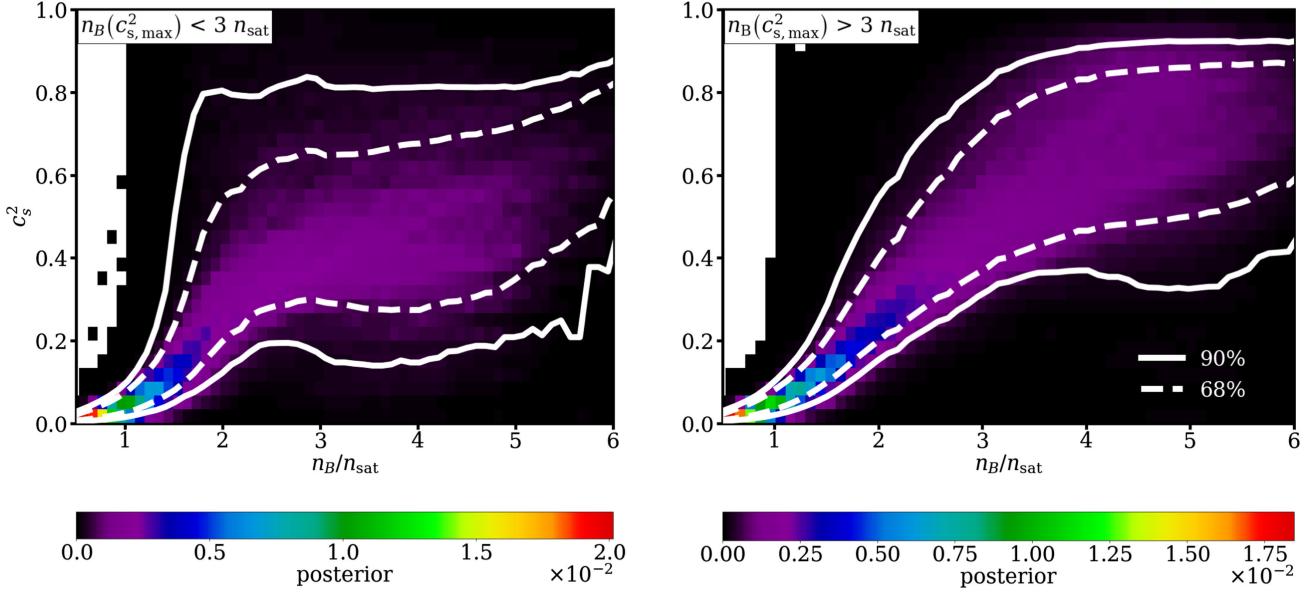


FIG. 9. EoS posteriors for the case when a global maximum in the speed of sound is present below (left) and above (right) $3n_{\text{sat}}$ for mGP EoS conditional on $n_B \leq n_B^{\text{max}}$ for each EoS. The EoS are represented by the speed of sound squared in units of c^2 as a function of baryon number density in units of n_{sat} . The posterior probability distributions are produced by binning the EoS by the speed of sound and number density, weighing each EoS by the corresponding likelihood, then normalizing the heights of the bins such that the sum of all bin heights is equal to 1. Also shown in the posterior plots are the 90% and 68% credible regions for the speed of sound squared at a given density for $0.5 \leq n_B \leq 6.0n_{\text{sat}}$. The posterior probability that the central density for a maximally massive star is greater than $\sim 6.0n_{\text{sat}}$ is negligible in both cases.

functional do not lead to significantly different predictions for the mass-radius relation or $c_s^2(n_B)$ given current constraints. On the other hand, introducing multiscale correlations via the mGP model had a significant impact on the posterior for the value of n_B at the maximum c_s^2 . What we learn from this is that both the benchmark GP and the mGP models can describe astronomical observations, while respecting symmetry energy and pQCD constraints. We can now ask if the data *prefer* one of the two models.

As discussed in Sec. IV A, the model evidence quantifies the level of support of the data for a given model, and the ratio between the evidence for two different models, the Bayes factor, quantifies if one of the models is preferred over the other by the data. Using Eq. (52), we separate the benchmark GP and the mGP samples and compute the evidence for each. We find a Bayes factor of⁹

$$K = \frac{\mathcal{E}_{\text{benchmark GP}}}{\mathcal{E}_{\text{mGP}}} = 1.480. \quad (60)$$

This value is not a significant deviation from unity, which means that current constraints do not favor either model. The physical interpretation is that multiscale correlations and nontrivial features in $c_s^2(n_B)$ are not ruled out by current constraints, but neither are they required.

Because of the broad range of features represented in the mGP samples, we further divide the mGP EoS into the different classes discussed in Sec. III A, resulting in four models: benchmark, mGP containing one spike, mGP containing a spike and a plateau, and mGP containing two plateaus. We present the log evidence for each of these classes separately in Table III. The log evidence across all classes of EoS is ~ 5 , so the log Bayes factor between any two classes is ~ 0 , and thus the Bayes factor between any

⁹In a previous paper [177], we reported a Bayes factor of $K = 1.126$, which was obtained using a normalization factor [N_m in Eq. (52)] that reflected the size of the subset of the priors for each model that passed the checks (using the notation introduced in Secs. III C and IV C, these sets are $\Phi_{\text{benchmark GP}} \cap \Phi_{\mathcal{V}}$ and $\Phi_{\text{mGP}} \cap \Phi_{\mathcal{V}}$). This choice essentially ensured a Bayes factor of ~ 1 because both priors have information about astrophysical constraints (Ref. [172] also pointed this out). Here, N_m corresponds to the full prior sample size (using the notation introduced in Secs. III C and IV C, these sets are $\Phi_{\text{benchmark GP}}$ and Φ_{mGP}).

TABLE III. Log evidence for the benchmark model and the three classes of modifications represented in the mGP model.

Model	$\log(\mathcal{E})$
Benchmark	-5.340
mGP, $(\widehat{s}\widehat{p})$	-5.452
mGP, $(\widehat{p}l_1, \widehat{s}\widehat{p})$	-5.572
mGP, (pl_1, pl_2)	-5.577

two classes is ~ 1 . We can conclude that the data do not favor a specific type of structure in $c_s^2(n_B)$.

We highlight that, because modifications are introduced on top of a benchmark model, there is overlap between the type of structure that modifications create in $c_s^2(n_B)$. For instance, a spike or plateau placed above the benchmark baseline may create a peak in $c_s^2(n_B)$, but a peak can also be achieved by a spike or plateau placed below the baseline. This nontrivial interaction between modifications and the baseline functional form of $c_s^2(n_B)$ is illustrated in Fig. 3. The solid blue line is an example of a global maximum in $c_s^2(n_B)$ created by the introduction of a plateau above the baseline value for c_s^2 . On the other hand, the dot-dashed blue line shows a maximum in $c_s^2(n_B)$ created by a plateau placed below the baseline. Yet another example is given by the dot-dashed green curve, where a plateau is introduced above the baseline which does not result in a maximum in $c_s^2(n_B)$, because the baseline value of c_s^2 at the end of the plateau is higher than the value of c_s^2 across the plateau. Thus, calculating the evidence for particular functional forms of $c_s^2(n_B)$ (e.g., monotonic, one peak, multiple local maxima) requires additional analysis which can be the subject of future work.

VI. CONCLUSIONS

Nuclear physics models with phase transitions and exotic degrees of freedom contain multiscale features that present as nontrivial structure in c_s^2 . In this work, we developed a new framework, which we named modified Gaussian processes, as a novel approach to producing functional forms of the EoS for the cold, catalyzed nuclear matter in neutron stars that each contain long-, medium-, and short-range correlations. These EoS can be generated with high computational efficiency and contain features that are indicative of the emergence of exotic degrees of freedom.

We compared our new mGP EoS model to a benchmark GP model that contains only long-range correlations (i.e., does not contain any short- or medium-range correlations), using a Bayesian analysis that incorporated astrophysical data, low-energy nuclear physics constraints, and input from pQCD calculations. From our Bayesian analysis, we found that both the benchmark GP and mGP models provide nearly equivalent results for the mass-radius, $c_s^2(n_B)$, and maximum central density posteriors. In fact, the Bayesian evidence for both models is nearly the same, leading to a Bayes factor of 1.5 between them. Thus, it is clear that we cannot rule out nontrivial features in the speed of sound from the data and also that these EoS are as valid as a smooth EoS, given current data.

Given the very similar posteriors for both the benchmark GP and the mGP models, one may wonder if there are any differences between the two. We found that the main difference between the two models is that mGP model allows for EoS that have bump in c_s^2 at low densities. In

fact, the posterior for the baryon density at which the maximum of c_s^2 occurs leads to a binomial distribution with peaks at $n_B \sim 2n_{\text{sat}}$ and $n_B \sim 5n_{\text{sat}}$. In contrast, the benchmark model only produces a definite peak at $n_B \sim 5n_{\text{sat}}$. Thus, we must conclude that the benchmark model is not adequately exploring the possibility of a bump in c_s^2 at low densities, due to long correlation lengths. We argue that it is important to explore the possibility of peaks around $n_B \sim 2n_{\text{sat}}$ because a global maximum at such densities is compatible with the onset of exotic phases in the core of neutron stars and, in this analysis, its existence is completely driven by astrophysical data.

We emphasize that interpreting the EoS posteriors shown in this work above $2n_{\text{sat}}$ warrants caution. Because the mean values of c_s^2 used for the benchmark GP and mGP increase with density (although the soft GP mean is $1/3$), both posteriors are centered around high values of c_s^2 (seen in Fig. 5). However, while the correlation length used for the benchmark GP suppresses low values of c_s^2 at high densities, the mGP model allows for low c_s^2 in the same regime, and therefore the posterior is dominated by the prior. This feature of the mGP model is not apparent in Fig. 5, but it can be seen clearly in Fig. 9. More importantly, in this work, we do not attempt to infer the behavior of the EoS near the maximum central density—in fact, our results largely suggest that the likelihood is uninformative in that regime. Instead, we test whether astrophysical data and theory input prefer a smooth EoS displaying correlations across large ranges in pressure or EoS that allow for multiscale correlations that are compatible with nuclear physics models.

Another question we explored is if the presence of sharp features in c_s^2 is potentially excluded due to pQCD constraints at high n_B . Similar to what was done in Ref. [171], we applied the constraints at the maximum central density that is peaked between $5 - 6n_{\text{sat}}$ for both models. We find that these pQCD constraints only have a small effect on our results, with no visible effect on our posteriors. Only 0.0083% of all EoS in our study were inconsistent with the pQCD constraints entirely, and only 3.5% were in some degree of tension. Tension can occur because of uncertainty in an undetermined scale that arises from the contribution of missing higher-order terms. One sets a range of values to that scale, and some of those values may exclude an EoS, whereas others may not. The results presented in this analysis are restricted to our choice of EoS prior and matching density. The impact of different EoS priors as well as different choices for the matching density were studied in Ref. [162]. The authors found that the impact of pQCD on the EoS inference was very sensitive to the presence of EoS that display a drastic softening immediately above the matching density followed by a strong stiffening. We plan to further explore the consequences of the pQCD constraints and EoS prior choices in a follow-up analysis.

Other approaches have been used to tackle similar questions as studied here, such as a deep neural network [117] or linear segments in c_s^2 [178]. It would be relevant to directly compare these different methods to our mGP model in future work to study their ability to reproduce specific features in c_s^2 from nuclear physics models. Additionally, Essick *et al.* [172] developed a new technique to extract features indicative of phase transitions from functional forms of the EoS. Rather than modifying samples from a benchmark model, Essick *et al.* sample from a mixture of GPs that contain a wide range of correlation lengths that are fixed for a given EoS, including short correlation lengths that lead to behavior in c_s^2 that mimics arbitrary phase transitions. Using this method, which we emphasize does not include multiscale correlations in the EoS, Essick *et al.* also find that current data are not yet constraining enough to rule in favor or against nontrivial features in c_s^2 , with the exception of very strong first-order phase transitions (latent energy per particle $\gtrsim 100$ MeV). Interestingly, they also conclude from realistic simulations of future data that a catalog of around 100 events would at best lead to a Bayes factor of $\sim 10:1$ in favor of a phase transition, even when the true EoS contains a strong phase transition. Their conclusions support the idea that astrophysical observations of neutron stars, currently and in the near future, will not be able to constrain short-range correlations in the EoS.

Overall, the results presented here suggest that current constraints are not enough to rule definitively in favor of or against phase transitions to exotic degrees of freedom in the core of neutron stars and that unambiguous signatures of structure in the EoS still require investigation. A clear ruling regarding the existence of exotic matter in the core of neutron stars will require more precise input from astronomical observations, laboratory measurements, and input from effective theories and QCD at high densities [179]. Fortunately, more data are anticipated from the NICER Collaboration both in terms of better statistics on existing measurements and radii from new neutron stars. Additionally, LIGO/Virgo/KAGRA's fourth observing run started in May 2023 with better sensitivity than during the third observing run, such that more neutron star mergers that will provide Λ constraints are anticipated [180] and the binary love relation may provide further insight into structure in c_s^2 [88]. Finally, more nuclear physics data are anticipated from the Facility for Rare Isotope Beams that will help constrain the low-density regime of the EoS and from low-energy heavy-ion collisions that will probe the large density, low-temperature region of the QCD phase diagram [181].

ACKNOWLEDGMENTS

We thank Reed Essick for a detailed reading of the original version of the manuscript and thorough feedback,

Katerina Chatzioannou for insightful discussions on Gaussian processes, and Michael Wentzel for reviewing Appendix. D. M. is supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE—1746047 and the Illinois Center for Advanced Studies of the Universe Graduate Fellowship. J. N. H., D. M., and N. Y. were supported in part by the National Science Foundation (NSF) within the framework of the MUSES collaboration, under Grant No. OAC2103680. J. N. H. acknowledges financial support from the US-DOE Nuclear Science Grants No. DESC0020633 and No. DE-SC0023861. M. C. M. was supported in part by NASA ADAP Grant No. 80NSSC21K0649. N. Y. was supported in part by NSF Grant No. PHYS-2207650. We also acknowledge support from the Illinois Campus Cluster, a computing resource that is operated by the Illinois Campus Cluster Program (ICCP) in conjunction with the National Center for Supercomputing Applications (NCSA), which is supported by funds from the University of Illinois at Urbana-Champaign.

APPENDIX: QUANTIFYING THE IMPACT OF SPIKES ON THE $p(\epsilon)$ PLANE

In this appendix, we will motivate the inclusion of spikes in $c_s^2(\log_{10} p)$ and discuss how those modifications translate into the $\epsilon - p$ plane. We will show that, under the assumptions of the mGP, these spikes can result in significant modifications to the benchmark EoS on the $\epsilon - p$ plane.

In our work, we define a grid that is regularly spaced in $\log_{10} p$ with $\Delta \log_{10} p = 0.043$. That is, $\Delta p(p_i) = 10^{p_i+0.043} - 10^{p_i}$ grows with p_i . Additionally, to recover the energy density, we compute

$$\epsilon_{i+1} = \epsilon_i + \frac{1}{c_{s,i}^2} (p_{i+1} - p_i) \quad (\text{A1})$$

$$\Delta \epsilon_i = \frac{1}{c_{s,i}^2} \Delta p(p_i). \quad (\text{A2})$$

Thus, if a modification is introduced at p_j such that $c_{s,*}^2(p_j) \geq 0.1$ [since c_s^2 can be arbitrarily small, we will focus on cases where $c_s^2 = \mathcal{O}(0.1-1)$ first], the change in energy density at point j is at least $\mathcal{O}(\Delta \epsilon_j) = \mathcal{O}(\Delta p_j)$ and at most $\mathcal{O}(\Delta \epsilon_j) = 10\mathcal{O}(\Delta p_j)$. We can derive some rough bounds on $\Delta \epsilon_j$ by taking some representative values of pressure; we choose $\log_{10} p_{j,\text{low}} = 33.5$ erg cm⁻³ (in most EoS realizations, this value is slightly above nuclear saturation density) and $\log_{10} p_{j,\text{high}} = 36$ erg cm⁻³ (in most EoS realizations, this is close to the maximum central pressure for a stable star). We then have

$$\Delta p_{\text{low}} = 10^{p_{j,\text{low}}+0.043} - 10^{p_{j,\text{low}}} = \mathcal{O}(10^{-1}) \text{ MeV/fm}^3, \quad (\text{A3})$$

$$\Delta p_{\text{high}} = 10^{p_{j,\text{high}}+0.043} - 10^{p_{j,\text{high}}} = \mathcal{O}(100) \text{ MeV/fm}^3, \quad (\text{A4})$$

implying that when $c_{s,*}^2(p_j) \geq 0.1$, $\mathcal{O}(10^{-1} \text{ MeV/fm}^3) \leq \Delta \epsilon_j \leq \mathcal{O}(1000 \text{ MeV/fm}^3)$. Considering the energy density itself $\epsilon = \mathcal{O}(100 \text{ MeV/fm}^3)$, spike-type modifications for which c_s^2 remains above 0.1 (again, we will discuss the case when $c_s^2 < 0.1$ next) lead to changes of order 0.1%–1000% in the energy density. Therefore, Δp is not necessarily small, and increases in the energy density from one grid point to the next can be large in our analysis.

Now, let us focus on how a change in a single point on the $\log_{10} p$ - c_s^2 plane compares to its benchmark value on the ϵ - p plane. That is, we will estimate how different a spike-modified $\epsilon_{j,*}$ is from its benchmark value ϵ_j . Dropping the j index (from here on, it is implied we are talking about the point where a spike is introduced), the change in energy density due to a modification $c_{s,*}^2$ is

$$\Delta \epsilon_* = \frac{1}{c_{s,*}^2} \Delta p, \quad (\text{A5})$$

such that the ratio between the benchmark value and the spike-modified value above is

$$\frac{\Delta \epsilon}{\Delta \epsilon_*} = \frac{\Delta p}{c_s^2} \frac{c_{s,*}^2}{\Delta p} = \frac{c_{s,*}^2}{c_s^2}. \quad (\text{A6})$$

We can then write

$$\frac{\Delta \epsilon}{\Delta \epsilon_*} = \mathcal{R} = \frac{c_{s,*}^2}{c_s^2}, \quad (\text{A7})$$

using \mathcal{R} to represent the ratio. Because of causality, if both $c_{s,*}^2 \geq 0.1$ and $c_s^2 \geq 0.1$, then $\mathcal{R} = \mathcal{O}(0.1\text{--}10)$, meaning that $\mathcal{O}(\Delta \epsilon_*)$ is at least 0.1 $\mathcal{O}(\Delta \epsilon)$ and at most 10 $\mathcal{O}(\Delta \epsilon)$. We already discussed above that $\Delta \epsilon$ can be as small as $\mathcal{O}(0.1)$ and as large as $\mathcal{O}(100) \text{ MeV/fm}^3$. This shows explicitly that a spike-type modification in c_s^2 can translate into a large modification in $p(\epsilon)$.

Now, let us relax the assumption that $c_s^2 \geq 0.1$ and $c_s^2 \geq 0.1$. From Eq. (A7), it is clear that if $c_{s,*}^2 \gg c_s^2$, $\mathcal{R} \rightarrow 0$, and $\Delta \epsilon \ll \Delta \epsilon_*$. Therefore, when a spike is introduced such that $c_s^2 \lesssim 0.1$, the change in the $p(\epsilon)$ compared to the benchmark EoS could span different orders of magnitude.

[1] G. Baym, T. Hatsuda, T. Kojo, P. D. Powell, Y. Song, and T. Takatsuka, *Rep. Prog. Phys.* **81**, 056902 (2018).

[2] S. Borsányi, Z. Fodor, J. N. Guenther, R. Kara, S. D. Katz, P. Parotto, A. Pásztor, C. Ratti, and K. K. Szabó, *Phys. Rev. Lett.* **126**, 232001 (2021).

[3] S. Borsanyi, J. N. Guenther, R. Kara, Z. Fodor, P. Parotto, A. Pasztor, C. Ratti, and K. K. Szabo, *Phys. Rev. D* **105**, 114504 (2022).

[4] J. W. Holt and N. Kaiser, *Phys. Rev. C* **95**, 034326 (2017).

[5] C. Drischler, J. W. Holt, and C. Wellenhofer, *Annu. Rev. Nucl. Part. Sci.* **71**, 403 (2021).

[6] S. Huth *et al.*, *Nature (London)* **606**, 276 (2022).

[7] G. Raaijmakers, S. K. Greif, K. Hebeler, T. Hinderer, S. Nissanke, A. Schwenk, T. E. Riley, A. L. Watts, J. M. Lattimer, and W. C. G. Ho, *Astrophys. J. Lett.* **918**, L29 (2021).

[8] F. Douchin and P. Haensel, *Astron. Astrophys.* **380**, 151 (2001).

[9] F. Gulminelli and A. R. Raduta, *Phys. Rev. C* **92**, 055803 (2015).

[10] X. Viñas, C. Gonzalez-Boquera, M. Centelles, C. Mondal, and L. M. Robledo, *Symmetry* **13**, 1613 (2021).

[11] M. G. Alford, L. Brodie, A. Haber, and I. Tews, *Phys. Rev. C* **106**, 055804 (2022).

[12] A. Kurkela, P. Romatschke, and A. Vuorinen, *Phys. Rev. D* **81**, 105021 (2010).

[13] A. Kurkela, E. S. Fraga, J. Schaffner-Bielich, and A. Vuorinen, *Astrophys. J.* **789**, 127 (2014).

[14] T. Gorda, A. Kurkela, R. Paatelainen, S. Säppi, and A. Vuorinen, *Phys. Rev. Lett.* **127**, 162003 (2021).

[15] I. Legred, K. Chatzioannou, R. Essick, S. Han, and P. Landry, *Phys. Rev. D* **104**, 063003 (2021).

[16] E. Annala, T. Gorda, J. Hirvonen, O. Komoltsev, A. Kurkela, J. Nättälä, and A. Vuorinen, *Nat. Commun.* **14**, 8451 (2023).

[17] E. Annala, T. Gorda, E. Katerini, A. Kurkela, J. Nättälä, V. Paschalidis, and A. Vuorinen, *Phys. Rev. X* **12**, 011058 (2022).

[18] E. Annala, T. Gorda, A. Kurkela, J. Nättälä, and A. Vuorinen, *Nat. Phys.* **16**, 907 (2020).

[19] M. Marczenko, L. McLerran, K. Redlich, and C. Sasaki, *Phys. Rev. C* **107**, 025802 (2023).

[20] M. Buballa *et al.*, *J. Phys. G* **41**, 123001 (2014).

[21] M. Ferreira, R. Câmara Pereira, and C. Providência, *Phys. Rev. D* **102**, 083030 (2020).

[22] P. Jakobus, A. Motornenko, R. O. Gomes, J. Steinheimer, and H. Stoecker, *Eur. Phys. J. C* **81**, 41 (2021).

[23] M. Dutra, O. Lourenço, and D. P. Menezes, *Phys. Rev. C* **93**, 025806 (2016); **94**, 049901(E) (2016).

[24] L. McLerran and S. Reddy, *Phys. Rev. Lett.* **122**, 122701 (2019).

[25] J. J. Li, A. Sedrakian, and M. Alford, *Phys. Rev. D* **101**, 063022 (2020).

[26] M. G. Alford and A. Sedrakian, *Phys. Rev. Lett.* **119**, 161104 (2017).

[27] A. Zacchi, M. Hanuske, and J. Schaffner-Bielich, *Phys. Rev. D* **93**, 065011 (2016).

[28] D. E. Alvarez-Castillo, D. B. Blaschke, A. G. Grunfeld, and V. P. Pagura, *Phys. Rev. D* **99**, 063010 (2019).

[29] A. Zacchi and J. Schaffner-Bielich, *Phys. Rev. D* **100**, 123024 (2019).

[30] T. Zhao and J. M. Lattimer, *Phys. Rev. D* **102**, 023021 (2020).

[31] T. Minamikawa, T. Kojo, and M. Harada, *Phys. Rev. C* **103**, 045205 (2021).

[32] D. C. Duarte, S. Hernandez-Ortiz, and K. S. Jeong, *Phys. Rev. C* **102**, 025203 (2020).

[33] S. Sen and L. Sivertsen, *Astrophys. J.* **915**, 109 (2021).

[34] R. D. Pisarski, *Phys. Rev. D* **103**, L071504 (2021).

[35] J. R. Stone, V. Dexheimer, P. A. M. Guichon, A. W. Thomas, and S. Typel, *Mon. Not. R. Astron. Soc.* **502**, 3476 (2021).

[36] J. I. Kapusta and T. Welle, *Phys. Rev. C* **104**, L012801 (2021).

[37] R. Somasundaram and J. Margueron, *Europhys. Lett.* **138**, 14002 (2022).

[38] A. Motornenko, J. Steinheimer, V. Vovchenko, S. Schramm, and H. Stoecker, *Phys. Rev. C* **101**, 034904 (2020).

[39] G. Baym, S. Furusawa, T. Hatsuda, T. Kojo, and H. Togashi, *Astrophys. J.* **885**, 42 (2019).

[40] G. Malfatti, M. G. Orsaria, I. F. Ranea-Sandoval, G. A. Contrera, and F. Weber, *Phys. Rev. D* **102**, 063008 (2020).

[41] H. Tan, T. Dore, V. Dexheimer, J. Noronha-Hostler, and N. Yunes, *Phys. Rev. D* **105**, 023018 (2022).

[42] D. C. Duarte, S. Hernandez-Ortiz, K. S. Jeong, and L. D. McLerran, *Phys. Rev. C* **107**, 065201 (2023).

[43] D. C. Duarte, S. Hernandez-Ortiz, K. S. Jeong, and L. D. McLerran, *Phys. Rev. D* **104**, L091901 (2021).

[44] D. C. Duarte, S. Hernandez-Ortiz, and K. S. Jeong, *Phys. Rev. C* **102**, 065202 (2020).

[45] T. Kojo, *Eur. Phys. J. A* **52**, 51 (2016).

[46] T. Dore, J. M. Karthein, I. Long, D. Mroczek, J. Noronha-Hostler, P. Parotto, C. Ratti, and Y. Yamauchi, *Phys. Rev. D* **106**, 094024 (2022).

[47] N. Cruz Camacho, R. Kumar, J. Peterson, V. Dexheimer, R. Haas, T. A. Manning, and J. Noronha-Hostler, arXiv: 2409.06837.

[48] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaboration), *Phys. Rev. Lett.* **119**, 161101 (2017).

[49] S. De, D. Finstad, J. M. Lattimer, D. A. Brown, E. Berger, and C. M. Biwer, *Phys. Rev. Lett.* **121**, 091102 (2018); **121**, 259902(E) (2018).

[50] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. Lett.* **121**, 161101 (2018).

[51] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Astrophys. J. Lett.* **892**, L3 (2020).

[52] Z. Arzoumanian, P. Baker, A. Brazier, S. Burke-Spoliar, S. Chamberlin, S. Chatterjee, B. Christy, J. M. Cordes, N. J. Cornish, F. Crawford *et al.*, *Astrophys. J.* **859**, 47 (2018).

[53] J. Antoniadis, P. C. Freire, N. Wex, T. M. Tauris, R. S. Lynch, M. H. Van Kerkwijk, M. Kramer, C. Bassa, V. S. Dhillon, T. Driebe *et al.*, *Science* **340**, 1233232 (2013).

[54] H. T. Cromartie, E. Fonseca, S. M. Ransom, P. B. Demorest, Z. Arzoumanian, H. Blumer, P. R. Brook, M. E. DeCesar, T. Dolch, J. A. Ellis *et al.*, *Nat. Astron.* **4**, 72 (2020).

[55] M. C. Miller *et al.*, *Astrophys. J. Lett.* **887**, L24 (2019).

[56] M. C. Miller *et al.*, *Astrophys. J. Lett.* **918**, L28 (2021).

[57] T. E. Riley *et al.*, *Astrophys. J. Lett.* **887**, L21 (2019).

[58] T. E. Riley *et al.*, *Astrophys. J. Lett.* **918**, L27 (2021).

[59] M. B. Tsang *et al.*, *Phys. Rev. C* **86**, 015803 (2012).

[60] B.-A. Li, P. G. Krastev, D.-H. Wen, and N.-B. Zhang, *Eur. Phys. J. A* **55**, 117 (2019).

[61] C. Drischler, R. J. Furnstahl, J. A. Melendez, and D. R. Phillips, *Phys. Rev. Lett.* **125**, 202702 (2020).

[62] C. Drischler, J. A. Melendez, R. J. Furnstahl, and D. R. Phillips, *Phys. Rev. C* **102**, 054315 (2020).

[63] D. Adhikari *et al.* (PREX Collaboration), *Phys. Rev. Lett.* **126**, 172502 (2021).

[64] B. T. Reed, F. J. Fattoyev, C. J. Horowitz, and J. Piekariewicz, *Phys. Rev. Lett.* **126**, 172503 (2021).

[65] T.-G. Yue, L.-W. Chen, Z. Zhang, and Y. Zhou, *Phys. Rev. Res.* **4**, L022054 (2022).

[66] C. E. Rhoades, Jr. and R. Ruffini, *Phys. Rev. Lett.* **32**, 324 (1974).

[67] T. Gorda, O. Komoltsev, and A. Kurkela, *Astrophys. J.* **950**, 107 (2023).

[68] M. C. Miller, C. Chirenti, and F. K. Lamb, 10.3847/1538-4357/ab4ef9 (2019).

[69] S. K. Greif, G. Raaijmakers, K. Hebeler, A. Schwenk, and A. L. Watts, *Mon. Not. R. Astron. Soc.* **485**, 5363 (2019).

[70] G. Raaijmakers, T. E. Riley, and A. L. Watts, *Mon. Not. R. Astron. Soc.* **478**, 2177 (2018).

[71] T. E. Riley, G. Raaijmakers, and A. L. Watts, *Mon. Not. R. Astron. Soc.* **478**, 1093 (2018).

[72] I. Legred, K. Chatzioannou, R. Essick, and P. Landry, *Phys. Rev. D* **105**, 043016 (2022).

[73] L. Lindblom, *Phys. Rev. D* **82**, 103011 (2010).

[74] L. Lindblom, *Phys. Rev. D* **97**, 123019 (2018).

[75] L. Lindblom, *Phys. Rev. D* **105**, 063031 (2022).

[76] J. S. Read, B. D. Lackey, B. J. Owen, and J. L. Friedman, *Phys. Rev. D* **79**, 124032 (2009).

[77] C. A. Raithel, F. Özel, and D. Psaltis, *Astrophys. J.* **844**, 156 (2017).

[78] E. Annala, T. Gorda, A. Kurkela, and A. Vuorinen, *Phys. Rev. Lett.* **120**, 172703 (2018).

[79] M. Marczenko, L. McLerran, K. Redlich, and C. Sasaki, *EPJ Web Conf.* **274**, 07014 (2022).

[80] K. Hebeler, J. M. Lattimer, C. J. Pethick, and A. Schwenk, *Astrophys. J.* **773**, 11 (2013).

[81] B. Margalit and B. D. Metzger, *Astrophys. J. Lett.* **850**, L19 (2017).

[82] F. Ozel, D. Psaltis, T. Guver, G. Baym, C. Heinke, and S. Guillot, *Astrophys. J.* **820**, 28 (2016).

[83] H. Tan, J. Noronha-Hostler, and N. Yunes, *Phys. Rev. Lett.* **125**, 261104 (2020).

[84] P. Landry and R. Essick, *Phys. Rev. D* **99**, 084049 (2019).

[85] M.-Z. Han, J.-L. Jiang, S.-P. Tang, and Y.-Z. Fan, *Astrophys. J.* **919**, 11 (2021).

[86] P. Landry, R. Essick, and K. Chatzioannou, *Phys. Rev. D* **101**, 123007 (2020).

[87] R. Essick, P. Landry, and D. E. Holz, *Phys. Rev. D* **101**, 063007 (2020).

[88] H. Tan, V. Dexheimer, J. Noronha-Hostler, and N. Yunes, *Phys. Rev. Lett.* **128**, 161101 (2022).

[89] C. K. Williams and C. E. Rasmussen, *Gaussian Processes for Machine Learning* (MIT Press, Cambridge, MA, 2006), Vol. 2.

[90] O. Komoltsev and A. Kurkela, *Phys. Rev. Lett.* **128**, 202701 (2022).

[91] R. Essick, I. Tews, P. Landry, and A. Schwenk, *Phys. Rev. Lett.* **127**, 192701 (2021).

[92] R. Essick, I. Tews, P. Landry, S. Reddy, and D. E. Holz, *Phys. Rev. C* **102**, 055803 (2020).

[93] H. Togashi, K. Nakazato, Y. Takehara, S. Yamamuro, H. Suzuki, and M. Takano, *Nucl. Phys.* **A961**, 78 (2017).

[94] S. Typel, M. Oertel, and T. Klähn, *Phys. Part. Nucl.* **46**, 633 (2015).

[95] M. Oertel, M. Hempel, T. Klähn, and S. Typel, *Rev. Mod. Phys.* **89**, 015007 (2017).

[96] G. Baym, C. Pethick, and P. Sutherland, *Astrophys. J.* **170**, 299 (1971).

[97] M. Baldo, I. Bombaci, and G. F. Burgio, *Astron. Astrophys.* **328**, 274 (1997), <https://adsabs.harvard.edu/full/1997A%26A...328.274B>.

[98] P. Danielewicz and J. Lee, *Nucl. Phys.* **A818**, 36 (2009).

[99] B. K. Agrawal, S. Shlomo, and V. K. Au, *Phys. Rev. C* **72**, 014310 (2005).

[100] P. G. Reinhard and H. Flocard, *Nucl. Phys.* **A584**, 467 (1995).

[101] B. K. Agrawal, S. Shlomo, and V. K. Au, *Phys. Rev. C* **68**, 031304 (2003).

[102] T. Gaitanos, M. Di Toro, S. Typel, V. Baran, C. Fuchs, V. Greco, and H. H. Wolter, *Nucl. Phys.* **A732**, 24 (2004).

[103] F. Grill, H. Pais, C. Providência, I. Vidaña, and S. S. Avancini, *Phys. Rev. C* **90**, 045803 (2014).

[104] N. K. Glendenning and S. A. Moszkowski, *Phys. Rev. Lett.* **67**, 2414 (1991).

[105] A. Akmal, V. R. Pandharipande, and D. G. Ravenhall, *Phys. Rev. C* **58**, 1804 (1998).

[106] M. Oertel, C. Providência, F. Gulminelli, and A. R. Raduta, *J. Phys. G* **42**, 075202 (2015).

[107] T. Schürhoff, S. Schramm, and V. Dexheimer, *Astrophys. J. Lett.* **724**, L74 (2010).

[108] V. Dexheimer, R. Negreiros, and S. Schramm, *Phys. Rev. C* **92**, 012801 (2015).

[109] V. Dexheimer and S. Schramm, *Astrophys. J.* **683**, 943 (2008).

[110] V. Dexheimer, *Pub. Astron. Soc. Aust.* **34**, E006 (2017).

[111] V. A. Dexheimer and S. Schramm, *Phys. Rev. C* **81**, 045201 (2010).

[112] V. Dexheimer, R. O. Gomes, T. Klähn, S. Han, and M. Salinas, *Phys. Rev. C* **103**, 025808 (2021).

[113] A. Clevinger, J. Corkish, K. Aryal, and V. Dexheimer, *Eur. Phys. J. A* **58**, 96 (2022).

[114] V. Dexheimer, R. de Oliveira Gomes, S. Schramm, and H. Pais, *J. Phys. G* **46**, 034002 (2019).

[115] V. Dexheimer, R. Negreiros, and S. Schramm, *Phys. Rev. C* **91**, 055808 (2015).

[116] O. Ivanytskyi and D. B. Blaschke, *Particles* **5**, 514 (2022).

[117] Y. Fujimoto, K. Fukushima, and K. Murase, *Phys. Rev. D* **101**, 054016 (2020).

[118] O. Ivanytskyi and D. Blaschke, *Phys. Rev. D* **105**, 114042 (2022).

[119] R. Kumar *et al.* (MUSES Collaboration), *Living Rev. Relativity* **27**, 3 (2024).

[120] R. Essick (private communication).

[121] Y. Fujimoto, T. Kojo, and L. D. McLerran, *Phys. Rev. Lett.* **132**, 112701 (2024).

[122] P. A. M. Guichon, K. Saito, E. N. Rodionov, and A. W. Thomas, *Nucl. Phys.* **A601**, 349 (1996).

[123] M. Shahrbaf, D. Blaschke, S. Typel, G. R. Farrar, and D. E. Alvarez-Castillo, *Phys. Rev. D* **105**, 103005 (2022).

[124] M. Hippert, E. S. Fraga, and J. Noronha, *Phys. Rev. D* **104**, 034011 (2021).

[125] I. Tews, J. Carlson, S. Gandolfi, and S. Reddy, *Astrophys. J.* **860**, 149 (2018).

[126] O. Ivanytskyi and D. Blaschke, *Eur. Phys. J. A* **58**, 152 (2022).

[127] A. Ayriyan, D. Blaschke, A. G. Grunfeld, D. Alvarez-Castillo, H. Grigorian, and V. Abgaryan, *Eur. Phys. J. A* **57**, 318 (2021).

[128] C. Mondal, M. Antonelli, F. Gulminelli, M. Mancini, J. Novak, and M. Oertel, *Mon. Not. R. Astron. Soc.* **524**, 3464 (2023).

[129] N.-B. Zhang and B.-A. Li, *Phys. Rev. C* **108**, 025803 (2023).

[130] S. Han and A. W. Steiner, *Phys. Rev. D* **99**, 083014 (2019).

[131] S. Benic, D. Blaschke, D. E. Alvarez-Castillo, T. Fischer, and S. Typel, *Astron. Astrophys.* **577**, A40 (2015).

[132] K. Yagi and N. Yunes, *Phys. Rep.* **681**, 1 (2017).

[133] T. Hinderer, *Astrophys. J.* **677**, 1216 (2008); **697**, 964(E) (2009).

[134] S. De, D. Finstad, J. M. Lattimer, D. A. Brown, E. Berger, and C. M. Biwer, *Phys. Rev. Lett.* **121**, 091102 (2018).

[135] G. Agazie *et al.* (NANOGrav Collaboration), *Astrophys. J. Lett.* **951**, L9 (2023).

[136] T. Gorda, O. Komoltsev, A. Kurkela, and A. Mazeliauskas, *J. High Energy Phys.* **06** (2023) 002.

[137] W. D. Myers and W. J. Swiatecki, *Nucl. Phys.* **81**, 1 (1966).

[138] W. D. Myers and W. J. Swiatecki, *Nucl. Phys.* **A601**, 141 (1996).

[139] C. Drischler, S. Han, J. M. Lattimer, M. Prakash, S. Reddy, and T. Zhao, *Phys. Rev. C* **103**, 045808 (2021).

[140] R. Essick, P. Landry, A. Schwenk, and I. Tews, *Phys. Rev. C* **104**, 065804 (2021).

[141] P. B. Demorest, T. Pennucci, S. Ransom, M. Roberts, and J. Hessels, *Nature (London)* **467**, 1081 (2010).

[142] E. Fonseca *et al.*, *Astrophys. J. Lett.* **915**, L12 (2021).

[143] K. H. Lo, M. C. Miller, S. Bhattacharyya, and F. K. Lamb, *Astrophys. J.* **776**, 19 (2013).

[144] M. C. Miller and F. K. Lamb, *Astrophys. J.* **808**, 31 (2015).

[145] M. C. Miller, *Astrophys. Space Sci. Libr.* **461**, 1 (2020).

[146] M. C. Miller and F. K. Lamb, *Eur. Phys. J. A* **52**, 63 (2016).

[147] T. Hinderer, *Astrophys. J.* **677**, 1216 (2008).

[148] T. Hinderer, *Astrophys. J.* **697**, 964 (2009).

[149] L. Wade, J. D. Creighton, E. Ochsner, B. D. Lackey, B. F. Farr, T. B. Litthenberg, and V. Raymond, *Phys. Rev. D* **89**, 103012 (2014).

[150] Properties of the binary neutron star merger GW170817, <https://dcc.ligo.org/LIGO-P1800061/public>, accessed: 8/28/2023.

[151] GW190425: Observation of a compact binary coalescence with total mass 3.4 Msun, <https://dcc.ligo.org/LIGO-P190425/public>, accessed: 8/28/2023.

[152] J.O. Andersen and M. Strickland, *Phys. Rev. D* **66**, 105001 (2002).

[153] T. Gorda, A. Kurkela, R. Paatelainen, S. Säppi, and A. Vuorinen, *Phys. Rev. D* **104**, 074015 (2021).

[154] T. Gorda, A. Kurkela, P. Romatschke, S. Säppi, and A. Vuorinen, *Phys. Rev. Lett.* **121**, 202701 (2018).

[155] T. Gorda and S. Säppi, *Phys. Rev. D* **105**, 114005 (2022).

[156] E. S. Fraga, A. Kurkela, and A. Vuorinen, *Astrophys. J. Lett.* **781**, L25 (2014).

[157] R. A. Schneider, [arXiv:hep-ph/0303104](https://arxiv.org/abs/hep-ph/0303104).

[158] A. Rebhan and P. Romatschke, *Phys. Rev. D* **68**, 025022 (2003).

[159] W. Cassing, *Nucl. Phys. A* **795**, 70 (2007).

[160] F.G. Gardim and F.M. Steffens, *Nucl. Phys. A* **825**, 222 (2009).

[161] A. Ipp and A. Rebhan, *J. High Energy Phys.* 06 (2003) 032.

[162] O. Komoltsev, R. Somasundaram, T. Gorda, A. Kurkela, J. Margueron, and I. Tews, *Phys. Rev. D* **109**, 094030 (2024).

[163] A. C. Sempowski, C. Drischler, R. J. Furnstahl, J. A. Melendez, and D. R. Phillips, [arXiv:2404.06323](https://arxiv.org/abs/2404.06323).

[164] C. Ecker, T. Gorda, A. Kurkela, and L. Rezzolla, [arXiv: 2403.03246](https://arxiv.org/abs/2403.03246).

[165] E. R. Most, L. J. Papenfort, V. Dexheimer, M. Hanuske, S. Schramm, H. Stöcker, and L. Rezzolla, *Phys. Rev. Lett.* **122**, 061101 (2019).

[166] R. W. Romani, D. Kandel, A. V. Filippenko, T. G. Brink, and W. Zheng, *Astrophys. J. Lett.* **908**, L46 (2021).

[167] R. W. Romani, D. Kandel, A. V. Filippenko, T. G. Brink, and W. Zheng, *Astrophys. J. Lett.* **934**, L17 (2022).

[168] V. Doroshenko, V. Suleimanov, G. Pühlhofer, and A. Santangelo, *Nat. Astron.* **6**, 1444 (2022).

[169] J. A. J. Alford and J. P. Halpern, [10.3847/1538-4357/acaf55](https://arxiv.org/abs/10.3847/1538-4357/acaf55) (2023).

[170] S. Altiparmak, C. Ecker, and L. Rezzolla, *Astrophys. J. Lett.* **939**, L34 (2022).

[171] R. Somasundaram, I. Tews, and J. Margueron, *Phys. Rev. C* **107**, L052801 (2023).

[172] R. Essick, I. Legred, K. Chatzioannou, S. Han, and P. Landry, *Phys. Rev. D* **108**, 043013 (2023).

[173] L. Brandes, W. Weise, and N. Kaiser, *Phys. Rev. D* **108**, 094014 (2023).

[174] P. T. H. Pang, L. Sivertsen, R. Somasundaram, T. Dietrich, S. Sen, I. Tews, M. W. Coughlin, and C. Van Den Broeck, *Phys. Rev. C* **109**, 025807 (2024).

[175] J. Takatsy, P. Kovacs, G. Wolf, and J. Schaffner-Bielich, *Phys. Rev. D* **108**, 043002 (2023).

[176] D. Oliinychenko, A. Sorensen, V. Koch, and L. McLerran, *Phys. Rev. C* **108**, 034908 (2023).

[177] D. Mroczeck, M. C. Miller, J. Noronha-Hostler, and N. Yunes, *J. Phys. Conf. Ser.* **2536**, 012006 (2023).

[178] R. Somasundaram, I. Tews, and J. Margueron, *Phys. Rev. C* **107**, 025801 (2023).

[179] A. Lovato *et al.*, [arXiv:2211.02224](https://arxiv.org/abs/2211.02224).

[180] Z. Carson, K. Chatzioannou, C.-J. Haster, K. Yagi, and N. Yunes, *Phys. Rev. D* **99**, 083016 (2019).

[181] A. Sorensen *et al.*, *Prog. Part. Nucl. Phys.* **134**, 104080 (2024).