

Equation of state at neutron-star densities and beyond from perturbative QCD

Oleg Komoltsev^{1,*}, Rahul Somasundaram^{2,3,†}, Tyler Gorda^{4,5,6,‡}, Alekski Kurkela^{1,§}, Jérôme Margueron^{7,8,||} and Ingo Tews^{2,¶}

¹*Faculty of Science and Technology, University of Stavanger, 4036 Stavanger, Norway*

²*Theoretical Division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*

³*Department of Physics, Syracuse University, Syracuse, New York 13244, USA*

⁴*Institut für Theoretische Physik, Goethe Universität, Max-von-Laue-Strasse 1, 60438 Frankfurt am Main, Germany*

⁵*Technische Universität Darmstadt, Department of Physics, 64289 Darmstadt, Germany*

⁶*ExtreMe Matter Institute EMMI, GSI Helmholtzzentrum für Schwerionenforschung GmbH, 64291 Darmstadt, Germany*

⁷*Institut de Physique des 2 infinis de Lyon, CNRS/IN2P3, Université de Lyon, Université Claude Bernard Lyon 1, F-69622 Villeurbanne Cedex, France*

⁸*International Research Laboratory on Nuclear Physics and Astrophysics, Michigan State University and CNRS, East Lansing, Michigan 48824, USA*



(Received 12 February 2024; accepted 30 April 2024; published 20 May 2024)

We explore the consequences of imposing robust thermodynamic constraints arising from perturbative quantum chromodynamics (QCD) when inferring the dense-matter equation-of-state (EOS). We find that the termination density, up to which the EOS modeling is performed in an inference setup, strongly affects the constraining power of the QCD input. This sensitivity in the constraining power arises from EOSs that have a specific form, with drastic softening immediately above the termination density followed by a strong stiffening. We also perform explicit modeling of the EOS down from perturbative-QCD densities to construct a new QCD likelihood function that incorporates additional perturbative-QCD calculations of the sound speed and is insensitive to the termination density, which we make publicly available.

DOI: [10.1103/PhysRevD.109.094030](https://doi.org/10.1103/PhysRevD.109.094030)

I. INTRODUCTION

Studying the impact of perturbative quantum chromodynamics (pQCD) calculations on the inference of the equation of state (EOS) of neutron stars (NSs) has recently become a promising area of research. The fundamental theory of the strong interaction, quantum chromodynamics (QCD) is nonperturbative at low energies but becomes perturbative at asymptotically high energies, which are reached either at high temperatures or at high densities, at about $20\text{--}40n_{\text{sat}}$ in baryon number density [1–5]. Here, $n_{\text{sat}} \approx 0.16 \text{ fm}^{-3}$ is the nuclear saturation density.

The cores of stable NSs explore densities that are most likely limited to about $5\text{--}8n_{\text{sat}}$, depending on the EOS which is realized in their cores [6]. This implies that the theoretical description of dense matter in stable NSs

requires the solution of QCD in a nonperturbative regime. At low densities, up to $1 - 2n_{\text{sat}}$, effective-field-theory (EFT) approaches, such as chiral EFT [7,8], are commonly employed to study the NS EOS [9–11]. Such calculations are typically extended to higher densities using model-agnostic extrapolations [9,12–30], conditioned with various astrophysical inputs. Some of these analyses have also included input from the fundamental QCD Lagrangian in the perturbative regime (henceforth referred to as “QCD input”) in the form of pQCD calculations at high densities [16–18,22,25,27–31]. These works have reported results that differ from those that do not account for this QCD input. In particular, the works including the QCD input have reported softening of the EOS at energy densities above $750 \text{ MeV}/\text{fm}^3$, for which there seems to be no evidence without this input.

The connection between the perturbative regime of QCD and NSs was first investigated in Ref. [16], by extending model-agnostic extrapolation functions of the EOS up to the densities where pQCD calculations become reliable and imposing the pQCD values as boundary conditions for these extrapolations. More recently, Komoltsev and Kurkela [32] have suggested a new method to link both

* oleg.komoltsev@uis.no

† rsomasun@syr.edu

‡ gorda@itp.uni-frankfurt.de

§ aleksi.kurkela@uis.no

|| j.margueron@ip2i.in2p3.fr

¶ itews@lanl.gov

density regimes, allowing them to “integrate backward,” i.e., to propagate the pQCD constraints to lower densities in a completely general, analytical, and model-independent manner using only the thermodynamic potential and the conditions of causality and mechanical stability. In practice, the method provides a necessary condition for an extrapolated EOS that has to be fulfilled at all densities below the validity range of pQCD calculations. This condition asserts that all the points of the EOS can be connected to the pQCD values by some causal and stable interpolation. The QCD input was shown to constrain the EOS down to densities of $\sim 2.2\text{--}2.5n_{\text{sat}}$, if otherwise only causality and mechanical stability are considered beyond chiral EFT. Within this construction, the only uncertainty in the constraints arises from the truncation error of the pQCD computation of the EOS at perturbative densities (for a comprehensive study on the pQCD truncation errors, see Ref. [33]). However, at these densities, astrophysical observations of NSs are naturally more informative.

While the construction suggested in Ref. [32] fully bypasses the need for modeling NSs (as the input to these constraints derives from high-energy particle physics, not from NSs), in order to interface the QCD input with these astrophysical observations several assumptions must be made. These include extrapolations involved in nuclear theory, assumptions involved when extracting NS properties from observational data, the prior assumptions of how EOS models are constructed, and up to what terminating number density n_{term} such modeling is performed. These assumptions and their interplay with the pQCD calculations are an important subject in discussions about the constraining power of astrophysical and QCD inputs.

The constraining power of pQCD calculations was studied in detail in two recent publications [29,34], and was later also investigated by Refs. [30,35–41]. These works have employed astrophysical input in the form of gravitational-wave data constraining the tidal deformability of NSs in binary mergers [42–44], radio measurements of NS masses [45–47], as well as simultaneous mass-radius measurements using x-ray observations [48–54]. They discussed whether the QCD input reduces the uncertainties of the inferred EOS when imposed on top of astrophysical inputs or whether the QCD and astrophysical inputs disfavor similar behaviors. In the latter case, the two inputs would corroborate one another, strengthening the conclusion that would have been obtained using only one.

In this work, we compare the analyses conducted in Refs. [29] and [34]. These studies arrived at different conclusions but used the same pQCD calculations and the same way of constraining the EOS [32], see Fig. 1. While both of these studies agreed that the QCD input impacted the EOS at the densities reached in massive NSs, Ref. [29] found that QCD input provides considerable additional information on top of other sources of data whereas Ref. [34] found this additional impact to be marginal.




FIG. 1. The construction of the pQCD constraint of Ref. [32]. The explicit form of the EOS is specified until $n = n_{\text{term}}$. At this point, the pQCD constraint considers all possibilities for connecting the EOS from n_{term} to n_{pQCD} , where pQCD calculations become reliable. The blue and red solid lines show EOSs that are compatible with pQCD at n_{term} , having a pQCD likelihood function of 1. However, for the blue EOS, the parameter space for the EOS at higher densities is large (indicated by the blue dashed lines) while the EOS above n_{term} are extremely constrained for the red EOS (indicated by red dashed lines). Finally, the black line corresponds to an EOS which cannot be connected to pQCD at n_{term} without violating causality or thermodynamic stability. For this EOS, the pQCD likelihood function is 0.

The conclusions of these two studies seem to be in contradiction, and the aim of this work is to investigate the origin of this apparent discrepancy.

In addition to this direct comparison between Refs. [29,34], we also investigate the EOS behavior beyond n_{term} that is required in order for the EOS to remain consistent with the QCD input to higher densities. We do so by generating ensembles of extensions of a given EOS beyond n_{term} , which we condition with the QCD input at $15n_{\text{sat}}$. In this context, we also introduce a concept of a pQCD tension index, which quantifies the degree to which a given EOS is in tension with the high-density pQCD constraint within the construction of Ref. [32], before being fully excluded.

Finally, we also introduce a method for modifying the QCD input itself by marginalizing over a prior model of high density EOSs constructed by extrapolating down from $n < n_{\text{pQCD}}$. Such marginalization has so far only been performed within parametric models [28,30,31] (cf. the discussion in Ref. [36]). Importantly, such a construction also permits us to condition the model with further pQCD information not used in the original construction. In particular, the speed of sound exhibits very small perturbative corrections to the free result down to densities as low as $\approx 25n_{\text{sat}}$. By conditioning the high-density prior with the speed of sound gives us a novel, more theory-informed QCD input at the price of increased model dependence.

The paper is organized as follows. In Sec. II, we review the construction of Ref. [32] and define the QCD input. We also introduce the pQCD tension index. In Sec. III, we discuss our computational setup, including the prior EOS models used below n_{term} and the astrophysical constraints that we use within our Bayesian analysis. In Sec. IV, we discuss our results. Finally, in Sec. V, we discuss our findings and limitations of our work.

II. REVIEW OF THE QCD INPUT

Following Refs. [29] and [32], we implement the QCD input as follows. Given the EOS at two different points, at a low density $\beta_L = (n_L, p_L, \mu_L)$ and at a high density $\beta_H = (n_H, p_H, \mu_H)$, we may ask whether these two points can be connected by *any* stable and causal EOS. Here n is the baryon density, p is the pressure, and μ is the baryon chemical potential. This question can be answered by exploring the bounds for EOS behavior between β_L and β_H .

The range of possible pressure differences between (μ_L, n_L) and (μ_H, n_H) are limited by the following two EOSs in this interval

$$n_{\text{pmin}}(\mu) = n_L \frac{\mu}{\mu_L}, \quad \mu_L \leq \mu < \mu_H \quad (1)$$

$$n_{\text{pmax}}(\mu) = n_H \frac{\mu}{\mu_H}, \quad \mu_L < \mu \leq \mu_H \quad (2)$$

namely those which have one segment of $c_s^2 = 1$ followed by a phase transition or vice versa, respectively. These EOSs bound the range of possible pressure differences to be

$$\Delta p_{\text{min}} = \frac{n_L}{2} \left(\frac{\mu_H^2}{\mu_L} - \mu_L \right), \quad (3)$$

$$\Delta p_{\text{max}} = \frac{n_H}{2} \left(\mu_H - \frac{\mu_L^2}{\mu_H} \right). \quad (4)$$

Given values for β_L and β_H , the pressure difference between these two points is fixed. Therefore the simple check (for detailed derivation see Refs. [29,32])

$$\Delta p = p_H - p_L \in [\Delta p_{\text{min}}, \Delta p_{\text{max}}] \quad (5)$$

determines whether the two points can be connected by a causal and stable EOS. This condition does not make use of any specific form of the EOS between the low- and high-density points but rather implies the existence of at least one EOS connecting the two limits. The condition on the pressure difference can be rewritten as

$$0 \leq I_{\text{pQCD}} \equiv \frac{\Delta p - \Delta p_{\text{min}}}{\Delta p_{\text{max}} - \Delta p_{\text{min}}} \leq 1, \quad (6)$$

where we introduce a new quantity I_{pQCD} , the pQCD tension index. If the index $I_{\text{pQCD}} \notin [0, 1]$ the EOS is excluded because there is no possible causal and stable interpolation. If the index obtains the values $I_{\text{pQCD}} = 0 (I_{\text{pQCD}} = 1)$, then the EOS is forced to follow $n_{\text{pmin}}(n_{\text{pmax}})$, respectively. In particular, if $I_{\text{pQCD}} = 1$, the EOS must contain a phase transition of size

$$\Delta n = n_H \frac{\mu_L}{\mu_H} - n_L. \quad (7)$$

For typical values of $n_L = n_{\text{term}} = n_{\text{TOV}}$, we obtain that the strength of phase transition reaches $\Delta n \sim 20n_{\text{sat}}$, followed by a $c_s^2 = 1$ segment of $\sim 10n_{\text{sat}}$. Note that if I_{pQCD} is not exactly $I_{\text{pQCD}} = 0$ or $I_{\text{pQCD}} = 1$, then both extreme EOSs of Eqs. (1) and (2) are excluded. For $0 < I_{\text{pQCD}} < 1$, the EOS is bounded in the n - μ - p -space by a nontrivial envelope whose projection in the plane of energy-density ε and pressure p is shown in Fig. 1 (for an explicit expression, see Ref. [32]).

Following Ref. [29], we adopt the scale-averaging interpretation of the pQCD uncertainties for β_H . We use a log-uniform distribution for the dimensionless renormalization scale $X = \bar{\Lambda}/[2(\mu_{\text{pQCD}}/3)]$ in the range $[1/2, 2]$, which can be written as $w(\log X) = \mathbf{1}_{[\log(1/2), \log(2)]}(\log X)$. Here $\bar{\Lambda}$ is the renormalization scale in the modified minimal subtraction scheme [1].¹ The (in)sensitivity to the in-principle arbitrary range of the variation of X has been studied in detail in [33].

The QCD likelihood function then reads

$$P(\text{QCD}|\text{EOS}) = \int d(\log X) w(\log X) \mathbf{1}_{[0,1]}(I_{\text{pQCD}}(X, \text{EOS})) \quad (8)$$

where $\mathbf{1}_{[0,1]}(I_{\text{pQCD}}(X, \text{EOS}))$ is an indicator function discarding EOSs with $I_{\text{pQCD}} \notin [0, 1]$ for a fixed value of X . In general, the index I_{pQCD} depends on all six quantities (n_L, p_L, μ_L) and (n_H, p_H, μ_H) , which for fixed values of $\mu_H = \mu_{\text{QCD}}$ reduces to $I_{\text{pQCD}}(X, \text{EOS}) = I_{\text{pQCD}}(X, n_L, p_L, \mu_L)$.

We note that the convergence of the pQCD series depends on the unphysical renormalization scale. In particular, the convergence is slowest for the smaller values of $X \approx 1/2$. This slow convergence renders the $X = 1/2$ results, which differ quantitatively from the higher values

¹Note that the factor $2/3$ in the definition of the dimensionless renormalization scale arises from the requirement that logarithmically enhanced terms in the perturbative expansion vanish at low orders. The 3 arises because the natural scale in pQCD is the quark number chemical potential $\mu_q \equiv \mu/3$ and the factor 2 is related to the coefficient of ε when regulating the integrals in $4 - 2\varepsilon$ dimensions in the minimal subtraction scheme.

of X , less trustworthy; see Ref. [33] for a detailed analysis of the convergence of the pQCD results.

In the following sections we apply the QCD likelihood function at different termination densities, meaning that (n_L, p_L, μ_L) corresponds to $(n_{\text{term}}, p_{\text{term}}, \mu_{\text{term}})$. For the high density input, $\beta_H = (n_{\text{pQCD}}, p_{\text{pQCD}}, \mu_{\text{pQCD}})$, we use the partial next-to-next-to-next-to-leading order pQCD computation of Ref. [4].

Note that in terms of the sound speed, the QCD input constrains the density n_{term} , which is given by a (logarithmic) integral of $c_s^{-2}(\mu)$ and the pressure p_{term} , which is the integral of this integral $n(\mu)$. Hence for a given NS EOS, the input is more sensitive to individual features in the speed of sound at lower densities and is less sensitive to those near n_{term} .

III. COMPUTATIONAL SETUP

In order to study the dependence of our results on the differences in the analyses between Refs. [29,34], we use two different models to generate our EOS sets: the parametric sound-speed model of Ref. [34] and the Gaussian-process (GP) model of the speed of sound of Ref. [29]. Note that similar models have also been employed by Refs. [10,20,27,39].

A. Prior Set 1: The sound-speed model

To generate the CSM set, we follow Ref. [34]. At the lowest densities, we fix the EOS to the crust model of Ref. [55] and do not explore uncertainties in the crust modeling. The crust EOS is then matched with the EOS of the core via a cubic spline in the c_s^2 versus n plane. For the core EOS, below a certain breakdown density n_b , we assume that the EOS can be accounted for by nucleonic degrees of freedom, which in this work is described by the meta-model of Refs. [56,57]. The meta-model is a density-functional approach, similar to the Skyrme approach, that allows one to incorporate nuclear-physics knowledge directly encoded in terms of the nuclear empirical parameters (NEPs). We vary the NEPs in a rather broad range in order to explore a wide range of nuclear density functionals that describe matter up to n_b . The parameter n_b is varied uniformly in the range $1-2n_{\text{sat}}$. Furthermore, we discard all models that do not satisfy constraints provided by the chiral EFT calculations of Ref. [58] up to nuclear saturation density.

Above n_b , we use the CSM and create a nonuniform grid in density between n_b and $25n_{\text{sat}}$ with 9 grid points, where the 9th point is fixed at $25n_{\text{sat}}$ and the other 8 points are drawn from a uniform probability distribution in density. At each grid point, the squared sound-speed is uniformly sampled between 0 and 1 in units of the speed of light. Then, the sound speeds at each grid point are connected by linear segments, generating a sound-speed profile $c_s^2(n)$ for the EOS. The density-dependent speed of sound $c_s^2(n)$ can

then be integrated to give the EOS, i.e., the pressure $p(n)$, the energy density $\varepsilon(n)$, and the baryon chemical potential $\mu(n)$, see Ref. [27] for more details. The CSM set used here consists of $\approx 16,000$ EOSs generated this way.

B. Prior Set 2: Gaussian-Process model

In the GP model, for densities below $n = 0.57n_{\text{sat}}$ we use the crust model by Baym, Pethick, and Sutherland [59]. Above the crust-core transition density, we perform a GP regression in an auxiliary variable $\phi(n) \equiv -\ln(1/c_s^2(n) - 1)$, choosing the prior on this variable from the multivariate Gaussian distribution,

$$\phi(n) \sim \mathcal{N}(-\ln(1/\bar{c}_s^2 - 1), K(n, n')), \quad (9)$$

with a Gaussian kernel $K(n, n') = \eta \exp(-(n - n')^2/2\ell^2)$ [29]. In our hierarchical model the hyperparameters of the GP η , ℓ , and \bar{c}_s^2 are themselves drawn from normal distributions,

$$\begin{aligned} \eta &\sim \mathcal{N}(1.25, 0.2^2), \\ \ell &\sim \mathcal{N}(1.0n_{\text{sat}}, (0.25n_{\text{sat}})^2), \\ \bar{c}_s^2 &\sim \mathcal{N}(0.5, 0.25^2). \end{aligned} \quad (10)$$

For $n \leq 1.1n_{\text{sat}}$, this prior is then conditioned on the chiral EFT results from Ref. [9], with the average and difference between the ‘‘soft’’ and ‘‘stiff’’ EOSs in that work taken to be the mean and 90% credible interval for the conditioning [29]. The final EOSs are then once again obtained by integrating $c_s^2(n)$. The GP used in this work consists of 120,000 EOSs generated in this manner.

C. Astrophysical input

The astrophysical inputs that we consider in our analysis are the following:

- (1) Radio-astronomy constraints: We consider the mass-measurements of PSR J0348 + 0432 with $M = 2.01 \pm 0.04M_{\odot}$ [45] and PSR J1624 – 2230 with $M = 1.928 \pm 0.017M_{\odot}$ [60]. We approximate these uncertainties as normal distributions, which is valid to good accuracy.
- (2) Gravitational-wave data from binary NS mergers: We consider the two-dimensional PDF for the tidal deformability and mass ratio at fixed chirp mass extracted from GW170817 by the LIGO-Virgo collaboration [44], using the PhenomPNRT waveform model with low-spin priors (see Fig. 12 of [44]).
- (3) NICER observations: We consider the simultaneous mass and radius measurement of PSR J0740 + 6620 using the NICER + XMM-Newton data. We have employed here the PDF from Ref. [53] (see their Fig. 1, right panel).

After conditioning our priors with the astrophysical input, we obtain a posterior sample of 4,000 EOSs for CSM and 46,000 EOSs for the GP model. These inputs are identical to those used in Ref. [29] (referred to as “Pulsars + $\tilde{\Lambda}$ ”) with one exception. We choose not to use the additional “Occam” factor related to the prior mass distribution of NSs, $(m_{\text{TOV}} - m_{\text{min}})^{-4}$ (one power per mass integral), which penalizes high-mass EOSs that are completely consistent with astrophysical data; see Ref. [61,62] for a discussion of this point. These inputs are the same as used in Ref. [34] with the exception of the NICER measurement of PSR J0030 + 0451. We have not included the latter in the present analysis as we found its impact on estimating the pQCD constraining power to be small.

Reference [29] analyzed the astrophysical as well as the QCD inputs using a Bayesian approach whereas Ref. [34] performed a hard cut analysis. We have checked that, as far as the treatment of the astrophysical data is concerned, this difference does not alter the conclusions of our paper regarding the impact of the QCD input. The role played by the different treatments of the QCD input is discussed in the next section.

IV. RESULTS

Next, we focus on the impact of the differences in methodology between Refs. [29] and [34]. The three major differences between the setups in these works can be summarized as follows:

- D1: Ref. [29] used a fully Bayesian approach and the predictions are analyzed in terms of posterior probability distributions whereas Ref. [34] performed a hard cut analysis and focused on the envelope associated with all EOS models for all studied quantities.
- D2: The functional basis of the EOS modeling in the two works was different: Ref. [29] employed the non-parametric GP model (Prior set 2), while Ref. [34] employed the parametric CSM (Prior set 1).
- D3: The two works made different choices for n_{term} : Ref. [29] chose $n_{\text{term}} = 10n_{\text{sat}}$ whereas Ref. [34] chose $n_{\text{term}} = n_{\text{TOV}} \approx 5-8n_{\text{sat}}$, which is different for each EOS. Hence, the works inferred the EOS in different density ranges.

To investigate point D3, we will discuss results for varying n_{term} expressed either in terms of multiples of the saturation density n_{sat} or in terms of the EOS-dependent central density of the maximum-mass neutron star, n_{TOV} .

A. Impact of the QCD input on the EOS

We show the QCD likelihood of Eq. (8) for the two EOS models, GP and CSM, in the upper panels of Fig. 2, where both models terminate at $n_{\text{term}} = n_{\text{TOV}}$. In particular, we show the sorted QCD likelihood for a large sample of EOSs drawn from the posterior distributions of the CSM and GP models after they have been conditioned on the astrophysical data. EOSs for which the likelihood is 1 are accepted for

all values of the pQCD renormalization scale X whereas EOSs that have a likelihood of 0 cannot be connected to pQCD with a causal and stable interpolation for any value of $X \in [1/2, 2]$. EOSs with likelihoods between the two limits can be connected to pQCD for some but not all values of X , see Fig. 3 for more details. Figure 2 clearly shows that the choice of EOS model has a weak impact on the QCD likelihood, illustrating that the difference D2 plays a minor role in the following analysis.

We observe that for $n_{\text{term}} = n_{\text{TOV}}$ and $\mu_H = 2.6$ GeV ($\mu_H = 2.4$ GeV) roughly 20% ($\sim 40\%$) of the EOSs cannot be connected to pQCD for all values of X and are therefore affected by the QCD input. For a very small fraction of EOS, the QCD likelihood drops below 20% but no EOS has a vanishing QCD likelihood for this choice of n_{term} . In contrast, the results are very different for $n_{\text{term}} = 10n_{\text{sat}}$, which was chosen in Ref. [29]. Now, roughly 50–60% of EOS have a vanishing QCD likelihood. This shows the strong impact of the choice of n_{term} on the pQCD constraining power and explains a large part of the difference in the conclusions of Refs. [29,34]. The strong differences between the two cases arise from the behavior the EOS must adopt beyond n_{TOV} to remain consistent with the QCD input at $10n_{\text{sat}}$. This is discussed further in Sec. IV C.

We quantify the impact of the QCD likelihood function as a function of n_{term} by computing the fraction of the marginalized likelihood (evidence) excluded by the QCD input

$$1 - \frac{\sum_i w_i^{\text{astro}} \cdot w_i^{\text{QCD}}}{\sum_i w_i^{\text{astro}}}, \quad (11)$$

where the index i denotes the EOS, and the summation is performed over all EOSs in the ensemble. The symbols w_i^{astro} and w_i^{QCD} correspond to the likelihood arising from the astrophysical and QCD inputs, respectively. In Ref. [63] this quantity was called the average consistency (see also Eq. (25) of Ref. [29]), and corresponds to one minus the average QCD likelihood in the posterior ensemble conditioned with only the astrophysical input, which is indicated by the hatched area. This quantity takes values between zero and one, with larger values indicating a larger impact of the QCD input.

This quantity is shown as the function of the termination density for the two EOS models in multiples of n_{sat} in the middle panels and relative to n_{TOV} in the bottom panels of Fig. 2. As the models are extended to higher n_{term} , the fraction of the astrophysical evidence cut by QCD increases rapidly; if the models are extended to $n = 1.2n_{\text{TOV}}$ this fraction increases to 0.4 (0.5 for $\mu_{\text{QCD}} = 2.4$ GeV). At a fixed density of $n_{\text{term}} = 6.5n_{\text{sat}}$, roughly corresponding to the approximate 68%-credible upper limit for the n_{TOV} , the average consistency is 0.4. However, we stress that the stiffest EOSs, i.e., those with the lowest QCD likelihoods,

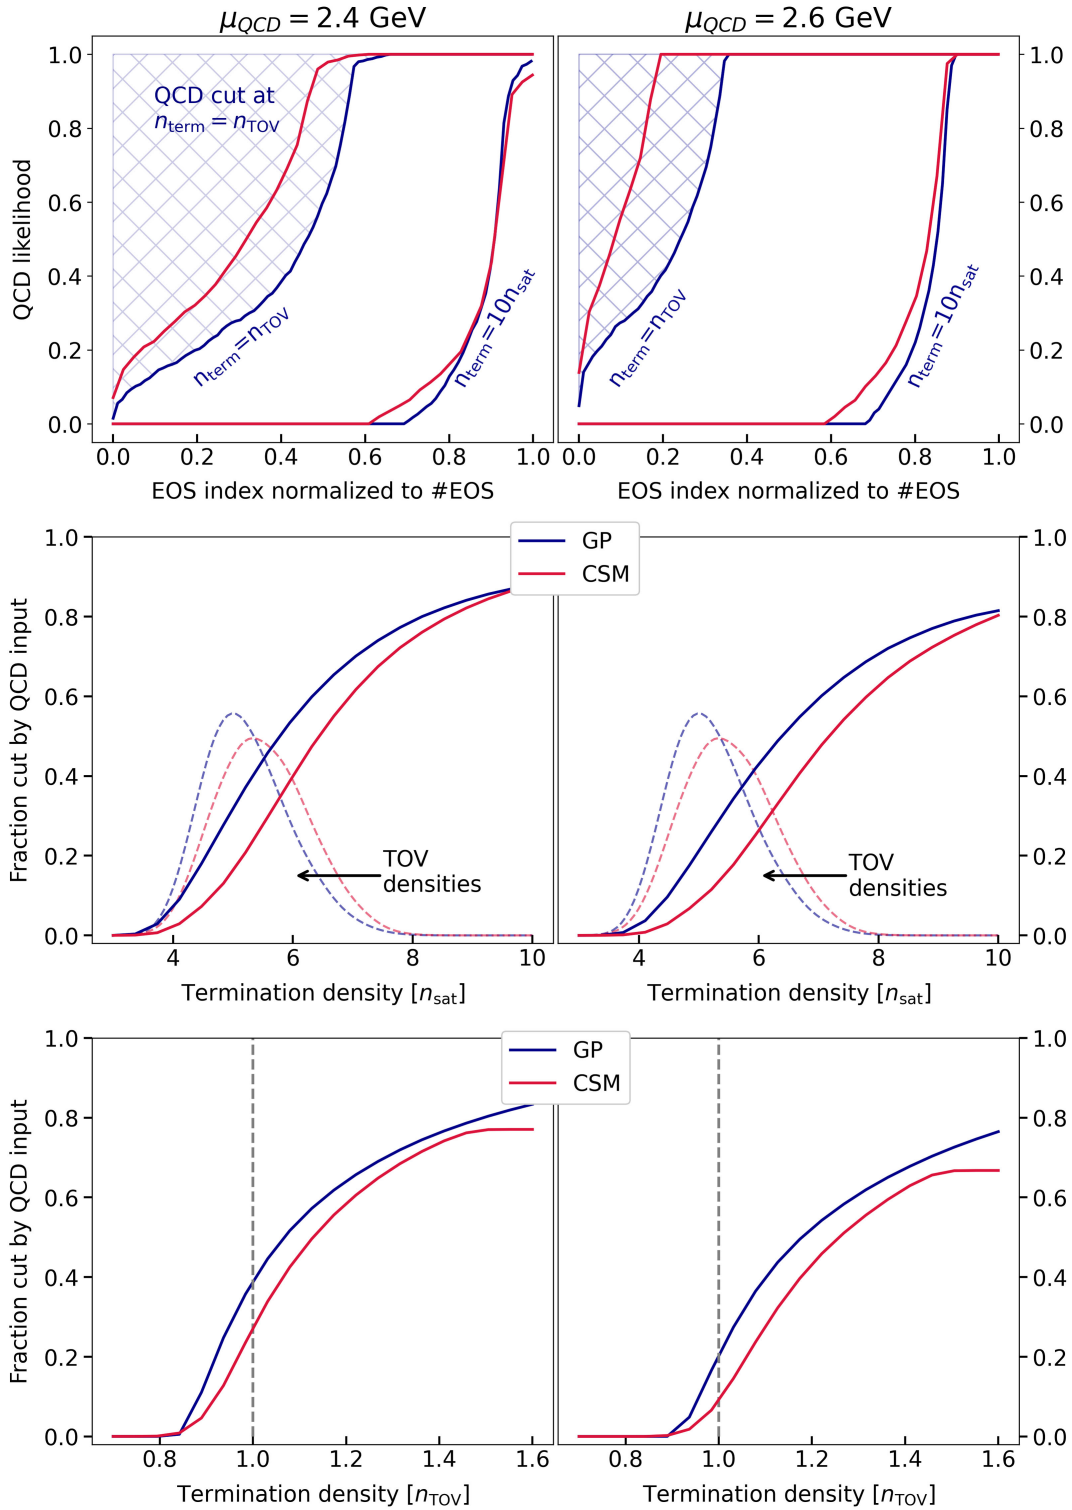


FIG. 2. Top: the values of the QCD likelihood function computed for a set of EOS draw from the posterior distribution. The posterior distributions for $\mu_{\text{pQCD}} = 2.4 \text{ GeV}$ (left panels) and $\mu_{\text{pQCD}} = 2.6 \text{ GeV}$ (right panels) have been conditioned on astrophysical data. The EOSs are ordered in increasing value of their QCD likelihood. The blue lines correspond to the GP, whereas the red lines correspond to the CSM. The likelihood function is displayed for two different termination densities, $n_{\text{term}} = n_{\text{TOV}}$ and $n_{\text{term}} = 10n_{\text{sat}}$, labeled in the figure. Middle: fraction of the astrophysical evidence (see Eq. (11)) that is cut by the QCD input as a function of the termination density. The dashed blue distribution shows the posterior for the GP whereas the dashed red lines show the same for the CSM. Bottom: fraction of the astrophysical evidence (see Eq. (11)) that is cut by the QCD input as a function of the termination density in units of the TOV density.

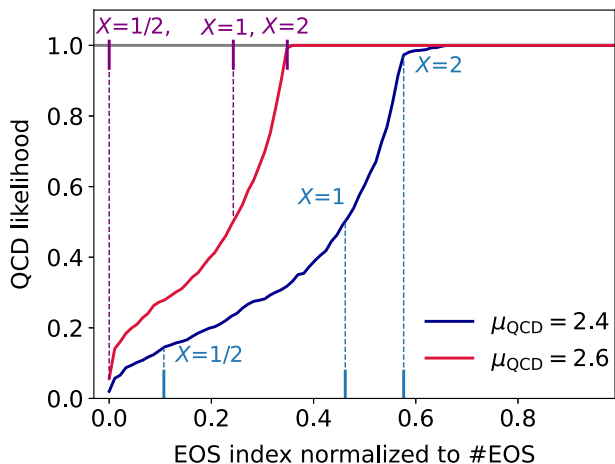


FIG. 3. The values of the QCD likelihood function at $n_{\text{term}} = n_{\text{TOV}}$ for a posterior sample of the GP model conditioned on astrophysical data. The EOSs are ordered according to their QCD likelihood. The blue line corresponds to $\mu_{\text{pQCD}} = 2.4$ GeV and a scale-averaged prescription for X in the range $[1/2, 2]$, whereas the red line corresponds to $\mu_{\text{pQCD}} = 2.6$ GeV. The colored ticks and dotted lines correspond to the likelihood function for fixed values of X ($X = 1/2, 1$ and 2), taking the form of a step function.

are in the lower percentile of the n_{TOV} range, and hence, for these EOS $n_{\text{TOV}} < 6.5n_{\text{sat}}$. These EOS are more strongly disfavored for larger n_{term} , when the EOS branch beyond n_{TOV} is probed. For higher termination densities, the average consistency levels out and takes values of ~ 0.8 at $n_{\text{term}} = 10n_{\text{sat}}$.

We, hence, find that the constraining power of the QCD input depends largely on the choice of the termination density n_{term} (D3). In fact, we find that a value of $n_{\text{term}} \approx 0.8n_{\text{TOV}}$ ($n_{\text{term}} \approx 4n_{\text{sat}}$ when applied at a fixed density)

constitutes a critical value, in the sense that it is at this point that the QCD input begins to show any effect at all beyond the astrophysical input. We note that the fact that this critical value of n_{term} is so close to n_{TOV} is a coincidence, in the sense that this critical value of n_{term} is set by a combination of many factors, including the value of n_{pQCD} at which pQCD is well converged and the constraining power of the astrophysical input (the latter applying only below n_{TOV}). In particular, if n_{pQCD} were significantly lower, then it is possible that this critical value of n_{term} would also be lower.

In order to investigate the impact of point D1 above, we show in Fig. 3 the detailed impact of X on the QCD likelihood. For different values of X , we indicate where the likelihood function becomes 1, i.e., all EOS to the right of these lines are accepted by the QCD input at $n_{\text{term}} = n_{\text{TOV}}$, while those on the left side are ruled out. Considering $\mu_{\text{pQCD}} = 2.6$ GeV as in Refs. [29,34], we find that nearly all EOS are accepted by $X = 1/2$, recovering the results of Ref. [34]. In contrast, roughly 25% of EOSs are rejected for $X = 1$ and 60% for $X = 2$. If μ_{pQCD} is decreased to $\mu_{\text{pQCD}} = 2.4$ GeV, about 10% of EOS are rejected already at $X = 1/2$. A Bayesian treatment of the scale parameter X as performed in Ref. [29] averages the impact of $X = 1/2$ leading to the findings of Ref. [29]. Hence we conclude that different statistical treatment of including pQCD calculations (D1) plays a secondary role in the differing conclusions of the two studies.

B. Impact of the QCD input on the ε - p values

We now analyze the effect of the QCD input on the allowed values of energy density and pressure. Figure 4 shows posterior credible regions of energy density ε and

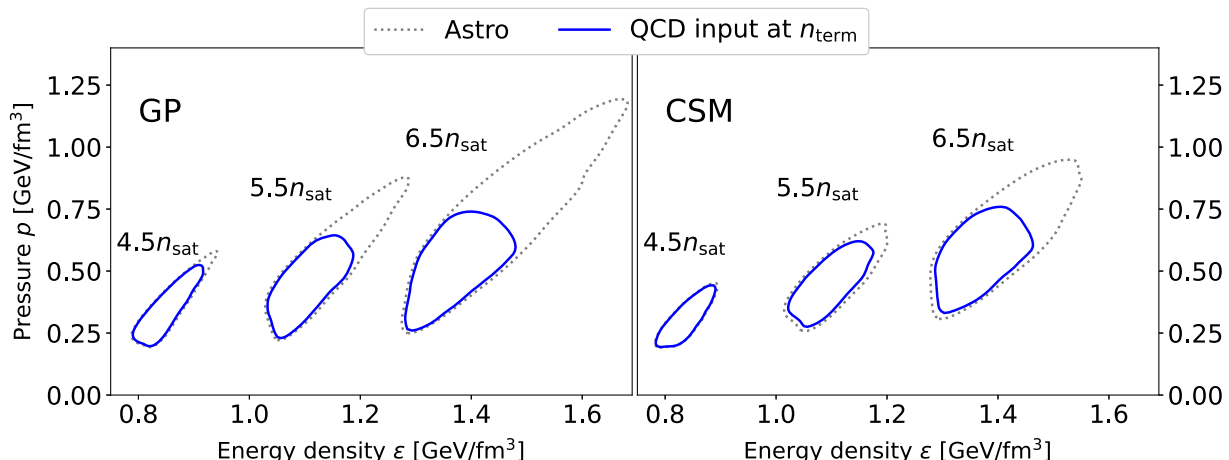


FIG. 4. Posterior probability density at fixed $n_{\text{term}} = 4.5, 5.5, 6.5n_{\text{sat}}$ of the energy density vs. pressure. The QCD input is enforced at n_{term} . The average TOV density for the GP posterior after enforcing astrophysical information is approximately $5.5n_{\text{sat}}$, while the maximum TOV density reaches approximately $8n_{\text{sat}}$ for the softest EOS. The areas correspond to 68% credible regions when considering only astrophysical information (dotted gray) and when also including the QCD input at n_{term} (solid blue). A value of $\mu_{\text{pQCD}} = 2.6$ GeV is used in this figure.

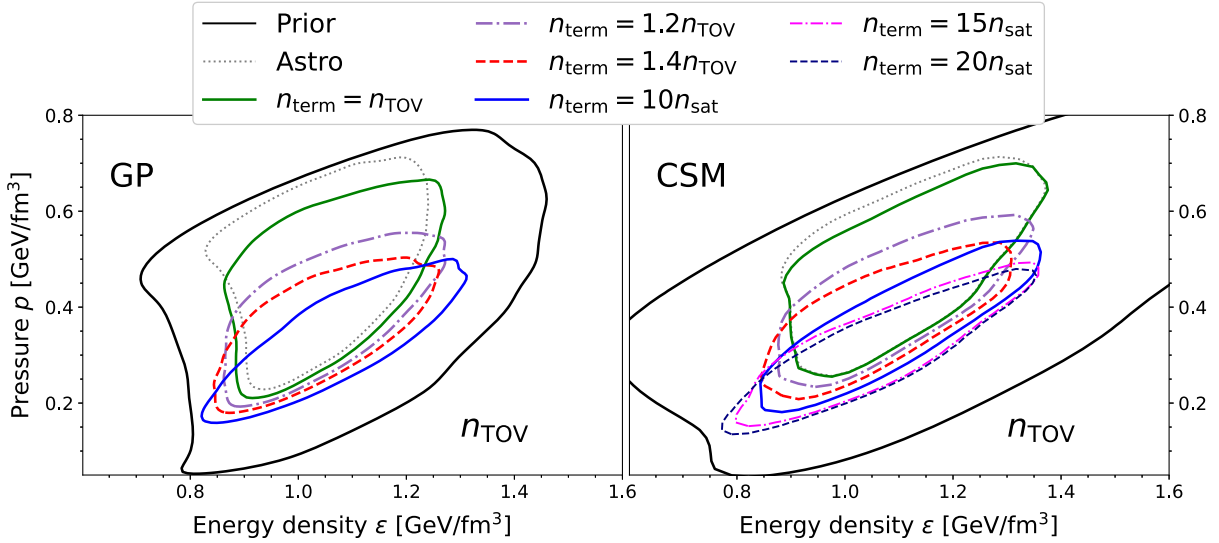


FIG. 5. Posterior probability density of the energy density vs. pressure at n_{TOV} for different termination densities. Each contour corresponds to 68% credible regions for the prior (black solid), considering only astrophysical input (dotted gray) and when also including QCD input (different colors correspond to different termination densities). The left plot shows results using the GP approach, while the right one shows results for the CSM approach. A value of $\mu_{\text{pQCD}} = 2.6$ GeV is used in this figure.

pressure p at various fixed densities in units of saturation density, incorporating either the astrophysical input alone or additionally also the QCD input. The EOS models are terminated at different densities $n_{\text{term}} \in \{4.5, 5.5, 6.5\}n_{\text{sat}}$, which correspond roughly to the range of n_{TOV} shown above in Fig. 2. We note that taking the termination density to be the density we are studying represents the maximally conservative approach to applying QCD constraint at a fixed density, as the QCD input is completely independent of the prior beyond n_{term} . As discussed in Ref. [29], the main effect of the QCD input is to disfavor the stiffest EOSs, visible in the figure. The effect of the input depends on the density; consistent with Fig. 2, at densities $\sim 6.5n_{\text{sat}}$ (corresponding to the 68%-credible upper limit on n_{TOV}) there is a significant reduction of allowed ϵ - p values, particularly at the highest p and ϵ . At lower densities the constraint is weaker as expected; at $5.5n_{\text{sat}}$, corresponding to the mean of n_{TOV} , the reduction is still significant, while taking $n_{\text{term}} = 4.5n_{\text{sat}}$ only shifts the posterior slightly at the highest pressures. This is in line with the prediction of Ref. [32] stating that QCD input ceases to give any constraint at $n < 2.2$ - $2.5n_{\text{sat}}$. From Fig. 4, we can also clearly see a difference between the two NS EOS models, namely that the GP model gives more weight to EOSs with larger values of p and ϵ than the CSM, visible in the posterior for the astrophysical input. This explains the slightly larger effect of the QCD input within the GP for the same n_{term} (point D2 above). However, after the QCD input is applied, the two models give quite similar posteriors.

We now move from studying the EOS at fixed densities to varying the termination density in terms of n_{TOV} . In Fig. 5, we show the posterior distribution of energy density

$\epsilon(n_{\text{TOV}})$ and pressure $p(n_{\text{TOV}})$ for $n_{\text{term}} \in \{1, 1.2, 1.4\}n_{\text{TOV}}$. The figure shows the 68%-credible regions for the prior distribution before introducing the astrophysical constraints as well as the regions after the EOSs have been conditioned on the astrophysical inputs but before the QCD input has been imposed. For both EOS models, the upper (left) edge of the credible region contour is nearly unchanged by the astrophysical input. This demonstrates that there is a region in the $\epsilon(n_{\text{TOV}})$ - $p(n_{\text{TOV}})$ -plane that is only marginally constrained by the present astrophysical input.

Similar to the ϵ - p -distribution at fixed density, the effect of the QCD input is to soften the EOS, disfavoring the highest pressures. We observe that for $n_{\text{term}} \in \{1.2, 1.4\}n_{\text{TOV}}$ the QCD input sets the upper left edge of the posterior distribution. This is also the case for $n_{\text{term}} = n_{\text{TOV}}$ in the GP model, while for the CSM at $n_{\text{term}} = n_{\text{TOV}}$ the impact is marginal. Consistent with Fig. 2, the quantitative impact is smallest for $n_{\text{term}} = n_{\text{TOV}}$ but the impact grows rapidly when n_{term} is increased and is qualitative similar to setting n_{term} to fixed density $10n_{\text{sat}}$, shown in Fig. 5. From both of these two figures, we observe that the choice of EOS prior (either GP or CSM) has only a minor impact on the effect of the QCD constraint. Finally, we also show contours corresponding to a fixed $n_{\text{term}} \in \{10, 15, 20\}n_{\text{sat}}$, where we show only CSM contours for the latter two values. We note that applying the QCD input for $n_{\text{term}} \in \{10, 15, 20\}n_{\text{sat}}$ yields very consistent credible regions.

C. Implications for the EOS beyond n_{term}

So far, we have discussed the impact the QCD input has on the EOS for $n < n_{\text{term}}$. In this subsection, we discuss how the QCD constraint affects the behavior of the EOS at

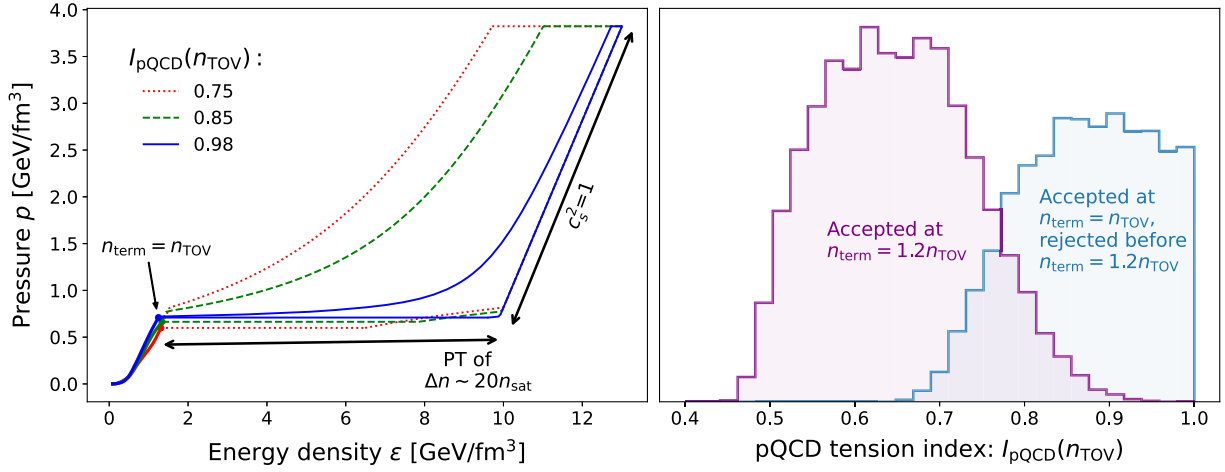


FIG. 6. Left: the QCD constraint shown in Fig. 1 for 3 different EOS with representative values of the pQCD tension index [see Eq. (6)]. Each EOS is terminated at $n_{\text{term}} = n_{\text{TOV}}$. The blue, dashed green, and dotted red areas envelop the region of EOS extensions that are consistent with the QCD input for a fixed $X = 1$ and $\mu_{\text{pQCD}} = 2.6$ GeV. The arrows illustrate the phase transition and the subsequent $c_s^2 = 1$ segment, which an EOS is forced to have when the pQCD tension index is close to 1. Right: the distribution of the pQCD tension index for EOSs that are accepted at $1.2n_{\text{TOV}}$ and, therefore, also at n_{TOV} (purple) as well as for EOS that are accepted at n_{TOV} , but rejected somewhere between n_{TOV} and $1.2n_{\text{TOV}}$ (blue). These distributions are normalized to the number of EOSs in each distribution. Most of the EOSs from the latter ensemble have a large pQCD tension index.

densities $n > n_{\text{term}}$. As shown in Fig. 1, the QCD input affects the available parameter space for the EOS in the p - ϵ plane beyond n_{term} , $n_{\text{term}} < n < n_H$. Some EOS have a large available parameter space while others are tightly constrained. One possible measure to gauge this freedom is the pQCD tension index I_{pQCD} defined before. The left panel of Fig. 6 shows the allowed space in the ϵ - p plane beyond n_{term} for three EOSs with $I_{\text{pQCD}}(n_{\text{TOV}}) = 0.75$, 0.85, and 0.98. The EOS with the highest I_{pQCD} must undergo a rapid phase-transition-like softening with $\Delta n \sim 20n_{\text{sat}}$, followed by a segment of $c_s^2 \approx 1$ to remain consistent with pQCD calculations. For $I_{\text{pQCD}} = 1$, this discontinuous behavior would be completely determined—see Fig. 7. Note that this phase transition is not responsible for the onset of the unstable NS branch, which instead is determined by general relativity within the density range modeled by the EOS. Instead, this phase transition takes place right above $n > n_{\text{TOV}}$. Therefore, the onset of this phase transition coinciding with n_{TOV} is accidental. If one wants to study the possibility of a phase transition determining the onset of the unstable branch, phase transitions must be included in the EOS models to a significant fraction, as in Ref. [64]. For $I_{\text{pQCD}} = 0.75$ a qualitatively similar softening is required but there is more freedom in the EOSs for $n > n_{\text{term}}$. By contrast, an EOS with $I_{\text{pQCD}} = 0.5$ is largely unconstrained above $n > n_{\text{TOV}}$.

In order to visualize the behavior above n_{term} explicitly, we have chosen representative EOSs that fulfill the QCD constraint at $n_{\text{term}} = n_{\text{TOV}}$ but not necessarily above. For each of these EOSs, we generate 5000 extensions between $n_{\text{TOV}} < n < 15n_{\text{sat}}$ using the GP approach. The extensions

are then conditioned with the QCD constraint at $15n_{\text{sat}}$ so that the behavior between $n_{\text{TOV}} < n < 15n_{\text{sat}}$ remains consistent with the QCD input. These extensions show how the EOS must evolve beyond n_{TOV} given its behavior below n_{TOV} . Figure 8 shows the extensions for representative EOSs with $I_{\text{pQCD}} = 0.5$, 0.75 and 0.85 (for $I_{\text{pQCD}} = 0.98$, our prior does not contain any valid samples). We have checked that these EOSs are indeed representative of EOSs with these pQCD tension indices. We find a difference in behavior above n_{TOV} based on the




FIG. 7. The behavior an EOS with tension index $I_{\text{pQCD}} = 1$ at n_{TOV} (blue line) has to follow above the termination density $n > n_{\text{term}}$ to match with pQCD for $X = 1$ (black line). The purple band corresponds to pQCD calculations of the speed of sound starting from $\mu_{\text{pQCD}} = 2.2$ GeV, with X varying in the range $[1/2, 2]$.




FIG. 8. Ensembles of possible extensions for three representative EOSs that are allowed at $n_{\text{term}} = n_{\text{TOV}}$. The EOSs correspond to different values for the pQCD tension index $I_{\text{pQCD}}(n_{\text{TOV}})$. The extensions are conditioned to fulfill the QCD constraint at $15n_{\text{sat}}$. In this figure, the coloring of each EOS is given by the value of $I_{\text{pQCD}}(15n_{\text{sat}})$.

value of the pQCD tension index at the TOV point $I_{\text{pQCD}}(n_{\text{TOV}})$. For the highest index of $I_{\text{pQCD}}(n_{\text{TOV}}) = 0.85$ shown in Fig. 8, we observe that the EOS must soften drastically above n_{TOV} , with $c_s^2(n) \leq 0.1$ all the way to $n = 15n_{\text{sat}}$. Furthermore, for this EOS, the value of the index at $15n_{\text{sat}}$ remains in the range $I_{\text{pQCD}}(15n_{\text{sat}}) \in [0.9, 1]$, indicating that the EOSs need to remain very soft. For the smaller value of the pQCD tension index $I_{\text{pQCD}}(n_{\text{TOV}}) = 0.75$, we similarly observe that the index at $15n_{\text{sat}}$ remains close to its value at n_{TOV} and the EOS must continue to remain soft. Finally, for the lowest $I_{\text{pQCD}}(n_{\text{TOV}}) = 0.5$ shown, the EOS extensions can

obtain significantly different values of the pQCD tension index at $15n_{\text{sat}}$.

We quantify the degree of softening as a function of pQCD tension index in Fig. 9, where for 100 EOSs in our posterior with fixed values of the pQCD tension index $I_{\text{pQCD}}(n_{\text{TOV}})$, we show the distribution of the average speed of sound of the extensions on the interval $n \in [8, 15]n_{\text{sat}}$. From this figure we see that the average c_s^2 for the extensions of EOSs with $I_{\text{pQCD}} = 0.85, 0.75$ are strongly skewed to small values, below $c_s^2 \lesssim 0.03$ and $c_s^2 \lesssim 0.11$, respectively (95% credence). By contrast, the extensions of EOSs with $I_{\text{pQCD}} = 0.5$ have a very large spread in the




FIG. 9. The distributions of the averaged sound speed in the density interval $[8, 15]n_{\text{sat}}$ for the possible EOS extensions of 100 different EOSs up to $n_{\text{term}} = n_{\text{TOV}}$. For each EOS, we generate 1000 extensions. This distribution is shown for fixed values of the pQCD tension index $I_{\text{pQCD}}(n_{\text{TOV}}) = 0.85, 0.75, 0.6, 0.5$.

average c_s^2 , with a shape that follows the prior distribution, indicating that there is no forced behavior in this case.

In the previous sections, we have seen that the constraining power of the QCD input increases as function of the termination density when going from $n_{\text{term}} = n_{\text{TOV}}$ to $n_{\text{term}} = 1.2n_{\text{TOV}}$; see, e.g., Fig. 2. In Fig. 6 we show the distribution of the pQCD tension indices $I_{\text{pQCD}}(n_{\text{TOV}})$ for EOSs that are allowed at $n_{\text{term}} = 1.2n_{\text{TOV}}$ and for EOS that are allowed at $n_{\text{term}} = n_{\text{TOV}}$ but rejected before $n_{\text{term}} = 1.2n_{\text{TOV}}$. The majority of the EOS that are excluded at higher n_{term} are EOSs which have large indices $I_{\text{pQCD}}(n_{\text{TOV}})$, forcing the EOS to undergo a rapid softening beyond n_{TOV} . This behavior cannot be excluded by the rigorous physics principles behind our implementation of the QCD constraint, but it would be interesting to understand what microphysical behavior could lead to such behavior.

D. Marginalization over EOS extensions beyond n_{term}

The termination of the explicit EOS models at n_{term} puts the EOS below and above n_{term} in an asymmetric role; below the termination density our prior contains a large number of nonextreme EOSs by construction. Above n_{term} , even the most extreme EOS extension of Fig. 7 is realized without any penalty. This has been criticized in Ref. [36], in which the authors point out that the construction of Ref. [32] is similar to performing a maximization, rather than a marginalization over the likelihoods of different possible extensions. As our final point in this work, we address the marginalization over a set of extensions of the EOS beyond n_{term} . We do so by constructing a model of possible high-density EOSs down from the pQCD EOS at $\mu_{\text{pQCD}} = 2.6$ GeV, distinct from the prior model of NS EOSs below n_{term} . Having an explicit model for EOSs also above n_{term} allows us to incorporate further theoretical

inputs. Since the pQCD speed of sound is well converged at high densities, see Fig. 7, we can further condition this model with the pQCD speed of sound. We may do so either only at the pQCD matching point $n_{\text{pQCD}} \approx 40n_{\text{sat}}$ (referred to below as “prior”), or over a larger range of densities where the speed of sound is under good theoretical control, $[25, 40]n_{\text{sat}}$ (“conditioned”).

For modeling the EOS above n_{term} , we define another hierarchical GP model anchored to the (scale-averaged) pQCD EOS defined with the following hyperparameters:

$$\begin{aligned} \ell &\sim \mathcal{U}(1n_{\text{sat}}, 20n_{\text{sat}}), \\ \eta &\sim \mathcal{N}(1.25, 0.25^2), \\ \bar{c}_s^2 &\sim \mathcal{N}(0.3, 0.3^2), \end{aligned} \quad (12)$$

see Fig. 10 (left). Negative values of η and \bar{c}_s^2 are removed. Because of the large interval in density between the typical n_{term} and $n_{\text{QCD}} \approx 40n_{\text{sat}}$, we choose a more general prior for the correlation length ℓ here than was done in Ref. [29]. We also choose the prior on \bar{c}_s^2 to have a mean centered near the conformal value of $1/3$, but allow it to vary broadly within the model. To be conservative with the conditioned GP, we take the standard deviation of the GP at a given $n \in [25, 40]n_{\text{sat}}$ to equal twice of that of the scale-averaged pQCD calculation.

The procedure for applying this high-density prior as a model for EOS extensions is as follows. First, we use either the low-density GP or CSM to model the NS EOS up to a specified termination density n_{term} . Then, for this n_{term} , we calculate a QCD likelihood function by constructing the posterior distribution for different ε - p values that arise from marginalizing over the new high-density GP model above n_{term} . Concretely, we calculate the distribution by performing a kernel-density estimation of the ε - p values of the high-density GP model at n_{term} . This posterior distribution is then reinterpreted as a new QCD likelihood function within the NS EOS inference at n_{term} that marginalizes over the hierarchical model above, as opposed to the construction from Ref. [32], which instead returns a weight of 1 for any ε - p point within the allowed region at n_{term} for a fixed X . More details of this procedure are given in Appendix.

In Fig. 10 we show in the left panel a small sample of EOSs drawn from these models of EOS extensions, showing a broad range of possible behaviors. The conditioning of the magenta GP is clearly visible. In the middle panel we show the QCD likelihood function resulting from the conditioned GP as applied at $n_{\text{term}} = 5.5n_{\text{sat}}$, compared to the constraints arising from Ref. [32] for the chiral EFT and pQCD constraints. This value of n_{term} approximately corresponds to the mean of the n_{TOV} for our ensembles using only the astrophysical input (see Fig. 2). We see that the likelihood function is mostly peaked in the interior of the $X = 1$ pQCD region. As has been pointed out already




FIG. 10. Left: a sample of EOSs drawn from the hierarchical models of EOS extensions of the perturbative EOS to lower densities. The dashed prior GP is anchored to the pQCD EOS at $\mu_{\text{pQCD}} = 2.6$ GeV, while the solid magenta GP is conditioned with the pQCD speed of sound on the interval $n \in [25, 40]n_{\text{sat}}$ (see main text). Middle: the resulting QCD likelihood function at $n_{\text{term}} = 5.5n_{\text{sat}}$ for the GP conditioned on the pQCD $c_s^2(n)$ for $n \in [25, 40]n_{\text{sat}}$. This termination density corresponds approximately to the mean value of n_{TOV} for our ensembles after applying the astrophysical input. Shown in blue and green are the constraints arising from the chiral EFT calculation of Ref. [9] and the $X = 1$ pQCD calculations of Ref. [4] alone, using the construction of Ref. [32]. Right: the posterior ϵ - p regions at n_{TOV} after marginalizing over the EOS extensions between n_{term} and $40n_{\text{sat}}$.

in Ref. [64], the impact of the QCD input within the inference of the NS EOS occurs primarily along the upper boundary of the pQCD region in Fig. 10, since the astrophysical constraints already reduce the allowed ϵ - p region in the neighborhood of the other extreme boundaries (cf. also the discussion of the upper edge of the ϵ - p regions in Fig. 5).

Turning to the rightmost panel of Fig. 10, we see the effect of the marginalization over the GP model of EOS extensions directly as the change in the allowed ϵ - p region at n_{TOV} . As expected from our observations of the central panel, we see a shift toward smaller p values in this figure. The prior GP shows a similar shift to the effect of using a termination density that is about 20%–40% larger without the marginalization (cf. the middle and right columns of Fig. 5), while the conditioned GP shows a much more pronounced softening, more comparable to the effect of the QCD input applied at $n_{\text{term}} = 10n_{\text{sat}}$ (cf. Fig. 5). In both cases, the quantitative effect is to shift the posterior pressure distribution within the most massive NSs to smaller values, with the conditioned GP shifted more strongly. That the upper limit for the conditioned GP is significantly below that of the prior demonstrates that the additional constraint is indeed due to the conditioning and not the assumed prior. This decreased upper bound is natural, as the pQCD sound speed used in the conditioning is inconsistent with the stiff high-density segment seen in Fig. 7, thus affecting the EOS at lower densities as well.

Lastly, we observe that the credible regions of ϵ - p are almost indistinguishable between $n_{\text{term}} = n_{\text{TOV}}$ and $10n_{\text{sat}}$, demonstrating independence of this parameter when using the conditioned GP extension.

V. DISCUSSION AND CONCLUSIONS

In this paper, we have studied the impact of the high-density QCD input on the modeling of the dense-matter

EOS, and explained the different findings presented in the literature. In particular, using the most general QCD input we found that the constraining power of the QCD input depends sensitively on the chosen termination density n_{term} . For the conservative choice of $n_{\text{term}} = n_{\text{TOV}}$, the impact of the QCD input is largely removed, while the constraining power rises fast with n_{term} . This difference in constraining power just beyond n_{TOV} arises from EOSs that undergo a change in behavior that is qualitatively different from the assumed priors. In particular, we have shown that most EOSs that are removed for $n_{\text{term}} > n_{\text{TOV}}$ have a large pQCD tension index, leading to an abrupt softening above n_{TOV} that is disfavored when the EOS priors employed here are constructed to higher densities. Hence, to remain in agreement with the QCD input, the EOS needs either to soften below n_{TOV} or it has to exhibit a drastic softening immediately above n_{TOV} that persists for an extended density range before stiffening again dramatically, see also Ref. [22]. This behavior is seen in the EOS extensions depicted in Fig. 8. We reiterate that the strong softening coinciding with n_{TOV} is accidental, as n_{TOV} is already set by the behavior of the EOS at lower densities. While it cannot be excluded by the generic thermodynamic considerations we employ, it would be interesting to see what microphysical behavior would be necessary to realize such behavior within dense matter. Note, however, that high speeds of sound above $n \gtrsim 25n_{\text{sat}}$ are inconsistent with pQCD calculations.

We note that different choices for n_{term} , $n_{\text{term}} = n_{\text{TOV}}$ and $n_{\text{term}} > n_{\text{TOV}}$, have their merits and answer slightly different questions: the choice of the termination density eventually reflects the phenomenological question at hand. On the one hand, the EOS is a fundamental prediction of QCD and exists independently of NSs at densities above n_{TOV} . The EOS above n_{TOV} is required in modeling the postmerger phase of binary NS mergers




FIG. 11. The impact of QCD input at the densities reached in the cores of maximally massive NS using different prescription for the EOS extension. Left: pressure as a function of energy density. For those prescriptions with $n_{\text{term}} = n_{\text{TOV}}$ the 68% credible intervals for the pressure are extended only up to lower bound of the 95% credible interval of the TOV energy densities. Above this point we cannot access the posterior distribution $P(p|\epsilon)$ but have access only to the conditional distribution $P(p|\epsilon, \epsilon < \epsilon_{\text{TOV}})$. Right: the corresponding posterior 68% credible region of ϵ - p at n_{TOV} .

(see Refs. [65,66]²) and sets a prediction for the quantity that may be verified in the future using, e.g., lattice field theory.³ In order to address these questions, the EOS must be modeled beyond the densities reached in stable NSs, and $n_{\text{term}} > n_{\text{TOV}}$ is required. On the other hand, one can be more conservative and use $n_{\text{term}} = n_{\text{TOV}}$ if one chooses to infer only the EOS along the stable NS branch. This choice avoids prior-dependent modeling of the unstable branch, allowing for more robust conclusions. Similarly, for an even more conservative modeling of NSs, one could imagine terminating the EOS model at $n_{\text{term}} = n(2.1M_{\odot})$ which is the mass of the heaviest observed neutron star. Whatever application one has in mind, it is crucial that the QCD input in the form of Ref. [29] be applied at the highest density that will be studied. It is inconsistent to apply the QCD input at n_{term} and then consider the model for even higher densities $n > n_{\text{term}}$.

Since the pQCD speed of sound is well converged to even lower densities than the pressure [4,5], we have additionally constructed a QCD likelihood function that incorporates this additional information. We have done so by constructing a further GP model down from the pQCD EOS calculation and conditioning it with the pQCD calculation of c_s^2 . This also addresses the points raised in Ref. [36], as this new likelihood function involves marginalizing over the new GP above n_{term} . We find that the conditioned GP gives results that are independent of $n_{\text{term}} \geq n_{\text{TOV}}$ (see Fig. 10) and agrees with results obtained for $n_{\text{term}} = 10n_{\text{sat}}$, see Fig. 11. The conditioned likelihood

function is made publicly available in Ref. [69]. We also see in Fig. 11 that the unconditioned GP leads to similar results as choosing $n_{\text{term}} = 1.2n_{\text{TOV}}$ or by restricting the pQCD tension index $I_{\text{pQCD}}(n_{\text{TOV}}) < 0.75$.

Finally, we discuss the limitations of our study. An important finding of this work is that the QCD input removes EOS that are very stiff around n_{TOV} . However, we allow for extreme stiffening of the EOS at low densities, around $1-2n_{\text{sat}}$, where EFT approaches can otherwise provide constraints on the EOS and typically suggest the EOS to be soft [10,70]. In this work, we have not varied the maximal density up to which we consider EFT input, which might bias our EOS sets to be stiffer than if we marginalized over the maximal EFT density [24,71]. We also have not marginalized over the scale μ_{pQCD} where we evaluate the pQCD values; marginalizing over lower values than used here would provide a more restrictive QCD input [33]. Furthermore, we only considered two EOS priors that model the EOS in terms of the speed of sound and are, hence, rather similar. We have also not considered astrophysical data from other x-ray sources such as quiescent low-mass x-ray binaries or photospheric-radius-expansion burst, which decrease the upper bound on the radii with respect to the NICER x-ray data [72]. Finally, for the purpose of marginalizing over EOS extensions above n_{term} , it would be interesting to investigate other prior choices for the EOS model, such as GP models with a density dependent correlation length.

ACKNOWLEDGMENTS

We thank Reed Essick, Joonas Hirvonen, Déborah Mroczek, Jorge Noronha, Jaki Noronha-Hostler, Krishna Rajagopal, and Rachel Steinhorst for useful discussions. This work has been supported in part by the Deutsche Forschungsgemeinschaft (DFG, German

²Note that Ref. [65] found that multimessenger signals from binary NS merges are not effective at constraining the EOS above n_{TOV} , rendering the QCD input the only source of information in this density range in the near future.

³On this point, see in particular two recent works in Refs. [67,68].

Research Foundation) project-ID 279384907–SFB 1245, by the State of Hesse within the Research Cluster ELEMENTS (projectID 500/10.006), and by the ERC Advanced Grant “JETSET: Launching, propagation and emission of relativistic jets from binary mergers and across mass scales” (Grant No. 884631) (T. G.). R. S. acknowledges support from the Nuclear Physics from Multi-Messenger Mergers (NP3M) Focused Research Hub which is funded by the National Science Foundation under Grant Number 21-16686, and by the Laboratory Directed Research and Development program of Los Alamos National Laboratory under Project No. 20220541ECR. J. M. is supported by the CNRS/IN2P3 NewMAC project, and are also grateful to PHAROS COST Action MP16214 and to the LABEX Lyon Institute of Origins (ANR-10-LABX-0066) of the Université de Lyon for its financial support within the program *Investissements d’Avenir* (ANR-11-IDEX-0007) of the French government operated by the National Research Agency (ANR). The work of I. T. was supported by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics, under Contract No. DE-AC52-06NA25396, by the Laboratory Directed Research and Development program of Los Alamos National Laboratory under Project No. 20230315ER, and by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Scientific Discovery through Advanced Computing (SciDAC) NUCLEI program.

Author contributions: O. K. performed most of the numerical analysis, with contributions from R. S. T. G. wrote the scripts for the marginalization and conditioning of the high-density GP. All authors contributed to the preparation and revision of the manuscript.

APPENDIX: MARGINALIZATION ABOVE n_{TOV}

In this appendix we present some details related to the marginalization procedure discussed in Sec. IV D and elaborate on why using the likelihood function discussed in the section corresponds to marginalization over the EOS prior above n_{TOV} .

We start by having an EOS prior in the whole density region between the low density and pQCD limits, described by some prior probability distribution $P(\eta_l, \eta_h)$ where the collective variables η_l and η_h correspond to the EOS

(or more specifically speed of sound) below and above an arbitrary density n_{term} .

In computing a posterior distribution of a physical quantity Q depending on the EOS, we are to marginalize over the EOS prior

$$P(Q) = \int d\eta_l d\eta_h P(Q|\eta_l, \eta_h) P(\eta_l, \eta_h) \quad (\text{A1})$$

Now lets make a simplifying assumption that the prior below and above n_{term} is not correlated (or is weakly correlated) such that

$$P(\eta_l, \eta_h) = P(\eta_l) P(\eta_h) \quad (\text{A2})$$

This is a good approximation beyond the correlation length of the GP.

Now, if the quantity Q does not depend on the high-density EOS η_h in any other way than through the end point p_L and ε_L at n_{term}

$$P(Q|\eta_l, \eta_h) = P(Q|\eta_l; p_L, \varepsilon_L), \quad (\text{A3})$$

as is the case in stable NSs if $n_{\text{term}} = n_{\text{TOV}}$, we may perform the integral over the η_h independent of the observable Q by marginalizing over p_L and ε_L

$$P(Q) = \int d\eta_l d\eta_h P(Q|\eta_l, \eta_h) P(\eta_l) P(\eta_h) \quad (\text{A4})$$

$$= \int d\eta_l d\eta_h dp_L d\varepsilon_L P(Q|\eta_l; p_L, \varepsilon_L) P(\eta_l) \quad (\text{A5})$$

$$\times P(p_L, \varepsilon_L|\eta_h) P(\eta_h) \quad (\text{A6})$$

$$= \int d\eta_l dp_L d\varepsilon_L P(Q|\eta_l; p_L, \varepsilon_L) P(\eta_l) w(p_L, \varepsilon_L), \quad (\text{A7})$$

where

$$w(p_L, \varepsilon_L) \equiv \int d\eta_h P(p_L, \varepsilon_L|\eta_h) P(\eta_h) \quad (\text{A8})$$

is the QCD likelihood function plotted in the center panel of Fig. 10.

- [1] A. Kurkela, P. Romatschke, and A. Vuorinen, Cold quark matter, *Phys. Rev. D* **81**, 105021 (2010).
 [2] A. Kurkela and A. Vuorinen, Cool quark matter, *Phys. Rev. Lett.* **117**, 042501 (2016).

- [3] T. Gorda, A. Kurkela, P. Romatschke, S. Säppi, and A. Vuorinen, Next-to-next-to-next-to-leading order pressure of cold quark matter: Leading logarithm, *Phys. Rev. Lett.* **121**, 202701 (2018).

- [4] T. Gorda, A. Kurkela, R. Paatelainen, S. Säppi, and A. Vuorinen, Soft interactions in cold quark matter, *Phys. Rev. Lett.* **127**, 162003 (2021).
- [5] T. Gorda, R. Paatelainen, S. Säppi, and K. Seppänen, Equation of state of cold quark matter to $O(\alpha_s^3 \ln \alpha_s)$, *Phys. Rev. Lett.* **131**, 181902 (2023).
- [6] J. Schaffner-Bielich, *Compact Star Physics* (Cambridge University Press, Cambridge, England, 2020).
- [7] E. Epelbaum, H.-W. Hammer, and U.-G. Meissner, Modern theory of nuclear forces, *Rev. Mod. Phys.* **81**, 1773 (2009).
- [8] R. Machleidt and D. R. Entem, Chiral effective field theory and nuclear forces, *Phys. Rep.* **503**, 1 (2011).
- [9] K. Hebeler, J. M. Lattimer, C. J. Pethick, and A. Schwenk, Equation of state and neutron star properties constrained by nuclear physics and observation, *Astrophys. J.* **773**, 11 (2013).
- [10] I. Tews, J. Carlson, S. Gandolfi, and S. Reddy, Constraining the speed of sound inside neutron stars with chiral effective field theory interactions and observations, *Astrophys. J.* **860**, 149 (2018).
- [11] C. Drischler, K. Hebeler, and A. Schwenk, Chiral interactions up to next-to-next-to-next-to-leading order and nuclear saturation, *Phys. Rev. Lett.* **122**, 042501 (2019).
- [12] J. S. Read, B. D. Lackey, B. J. Owen, and J. L. Friedman, Constraints on a phenomenologically parameterized neutron-star equation of state, *Phys. Rev. D* **79**, 124032 (2009).
- [13] L. Lindblom, Spectral representations of neutron-star equations of state, *Phys. Rev. D* **82**, 103011 (2010).
- [14] L. Lindblom and N. M. Indik, A spectral approach to the relativistic inverse stellar structure problem, *Phys. Rev. D* **86**, 084003 (2012).
- [15] L. Lindblom and N. M. Indik, Spectral approach to the relativistic inverse stellar structure problem II, *Phys. Rev. D* **89**, 064003 (2014); *Phys. Rev. D* **93**, 129903 (2016).
- [16] A. Kurkela, E. S. Fraga, J. Schaffner-Bielich, and A. Vuorinen, Constraining neutron star matter with quantum chromodynamics, *Astrophys. J.* **789**, 127 (2014).
- [17] E. Annala, T. Gorda, A. Kurkela, and A. Vuorinen, Gravitational-wave constraints on the neutron-star-matter equation of state, *Phys. Rev. Lett.* **120**, 172703 (2018).
- [18] E. R. Most, L. R. Weih, L. Rezzolla, and J. Schaffner-Bielich, New constraints on radii and tidal deformabilities of neutron stars from GW170817, *Phys. Rev. Lett.* **120**, 261103 (2018).
- [19] I. Tews, J. Margueron, and S. Reddy, Critical examination of constraints on the equation of state of dense matter obtained from GW170817, *Phys. Rev. C* **98**, 045804 (2018).
- [20] P. Landry and R. Essick, Nonparametric inference of the neutron star equation of state from gravitational wave observations, *Phys. Rev. D* **99**, 084049 (2019).
- [21] S. K. Greif, G. Raaijmakers, K. Hebeler, A. Schwenk, and A. L. Watts, Equation of state sensitivities when inferring neutron star and dense matter properties, *Mon. Not. R. Astron. Soc.* **485**, 5363 (2019).
- [22] E. Annala, T. Gorda, A. Kurkela, J. Nättilä, and A. Vuorinen, Evidence for quark-matter cores in massive neutron stars, *Nat. Phys.* **16**, 907 (2020).
- [23] G. Raaijmakers, S. K. Greif, T. E. Riley, T. Hinderer, K. Hebeler, A. Schwenk, A. L. Watts, S. Nissanke, S. Guillot, J. M. Lattimer, and R. M. Ludlam, Constraining the dense matter equation of state with joint analysis of NICER and LIGO/Virgo measurements, *Astrophys. J. Lett.* **893**, L21 (2020).
- [24] R. Essick, I. Tews, P. Landry, S. Reddy, and D. E. Holz, Direct astrophysical tests of chiral effective field theory at supranuclear densities, *Phys. Rev. C* **102**, 055803 (2020).
- [25] E. Annala, T. Gorda, E. Katerini, A. Kurkela, J. Nättilä, V. Paschalidis, and A. Vuorinen, Multimessenger constraints for ultradense matter, *Phys. Rev. X* **12**, 011058 (2022).
- [26] S. Altiparmak, C. Ecker, and L. Rezzolla, On the sound speed in neutron stars, *Astrophys. J. Lett.* **939**, L34 (2022).
- [27] R. Somasundaram, I. Tews, and J. Margueron, Investigating signatures of phase transitions in neutron-star cores, *Phys. Rev. C* **107**, 025801 (2023).
- [28] J.-L. Jiang, C. Ecker, and L. Rezzolla, Bayesian analysis of neutron-star properties with parameterized equations of state: The role of the likelihood functions, *Astrophys. J.* **949**, 11 (2023).
- [29] T. Gorda, O. Komoltsev, and A. Kurkela, Ab-initio QCD calculations impact the inference of the neutron-star-matter equation of state, *Astrophys. J.* **950**, 107 (2023).
- [30] E. Annala, T. Gorda, J. Hirvonen, O. Komoltsev, A. Kurkela, J. Nättilä, and A. Vuorinen, Strongly interacting matter exhibits deconfined behavior in massive neutron stars, *Nat. Commun.* **14**, 8451 (2023).
- [31] M.-Z. Han, Y.-J. Huang, S.-P. Tang, and Y.-Z. Fan, Plausible presence of new state in neutron stars with masses above 0.98MTOV, *Sci. Bull.* **68**, 913 (2023).
- [32] O. Komoltsev and A. Kurkela, How perturbative QCD constrains the equation of state at neutron-star densities, *Phys. Rev. Lett.* **128**, 202701 (2022).
- [33] T. Gorda, O. Komoltsev, A. Kurkela, and A. Mazeliauskas, Bayesian uncertainty quantification of perturbative QCD input to the neutron-star equation of state, *J. High Energy Phys.* **06** (2023) 002.
- [34] R. Somasundaram, I. Tews, and J. Margueron, Perturbative QCD and the neutron star equation of state, *Phys. Rev. C* **107**, L052801 (2023).
- [35] C. Ecker and L. Rezzolla, Impact of large-mass constraints on the properties of neutron stars, *Mon. Not. R. Astron. Soc.* **519**, 2615 (2022).
- [36] R. Essick, I. Legred, K. Chatzioannou, S. Han, and P. Landry, Phase transition phenomenology with nonparametric representations of the neutron star equation of state, *Phys. Rev. D* **108**, 043013 (2023).
- [37] L. Brandes, W. Weise, and N. Kaiser, Evidence against a strong first-order phase transition in neutron star cores: Impact of new data, *Phys. Rev. D* **108**, 094014 (2023).
- [38] C. Musolino, C. Ecker, and L. Rezzolla, On the maximum mass and oblateness of rotating neutron stars with generic equations of state, *Astrophys. J.* **962**, 61 (2024).
- [39] D. Mroczek, M. C. Miller, J. Noronha-Hostler, and N. Yunes, Nontrivial features in the speed of sound inside neutron stars, [arXiv:2309.02345](https://arxiv.org/abs/2309.02345).
- [40] Y.-Z. Fan, M.-Z. Han, J.-L. Jiang, D.-S. Shao, and S.-P. Tang, Maximum gravitational mass $M_{\text{TOV}} = 2.25^{+0.08}_{-0.07} M_{\odot}$ inferred at about 3% precision with multimessenger data of neutron stars, *Phys. Rev. D* **109**, 043052 (2024).
- [41] S.-P. Tang, M.-Z. Han, Y.-J. Huang, Y.-Z. Fan, and D.-M. Wei, Mass and radius of the most massive neutron star:

- The probe of the equation of state and perturbative QCD, *Phys. Rev. D* **109**, 083037 (2024).
- [42] B. P. Abbott, R. Abbott, T. D. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), GW170817: Observation of gravitational waves from a binary neutron star inspiral, *Phys. Rev. Lett.* **119**, 161101 (2017).
- [43] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese *et al.* (The LIGO Scientific and the Virgo Collaborations), GW170817: Measurements of neutron star radii and equation of state, *Phys. Rev. Lett.* **121**, 161101 (2018).
- [44] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese *et al.* (LIGO Scientific and Virgo Collaborations), Properties of the binary neutron star merger GW170817, *Phys. Rev. X* **9**, 011001 (2019).
- [45] J. Antoniadis *et al.*, A massive pulsar in a compact relativistic binary, *Science* **340**, 1233232 (2013).
- [46] H. T. Cromartie *et al.* (NANOGrav Collaboration), Relativistic Shapiro delay measurements of an extremely massive millisecond pulsar, *Nat. Astron.* **4**, 72 (2019).
- [47] E. Fonseca *et al.*, Refined mass and geometric measurements of the high-mass PSR J0740 + 6620, *Astrophys. J. Lett.* **915**, L12 (2021).
- [48] A. W. Steiner, C. O. Heinke, S. Bogdanov, C. Li, W. C. G. Ho, A. Bahramian, and S. Han, Constraining the mass and radius of neutron stars in globular clusters, *Mon. Not. R. Astron. Soc.* **476**, 421 (2018).
- [49] J. Nättilä, M. C. Miller, A. W. Steiner, J. J. E. Kajava, V. F. Suleimanov, and J. Poutanen, Neutron star mass and radius measurements from atmospheric model fits to x-ray burst cooling tail spectra, *Astron. Astrophys.* **608**, A31 (2017).
- [50] A. W. Shaw, C. O. Heinke, A. W. Steiner, S. Campana, H. N. Cohn, W. C. G. Ho, P. M. Lugger, and M. Servillat, The radius of the quiescent neutron star in the globular cluster M13, *Mon. Not. R. Astron. Soc.* **476**, 4713 (2018).
- [51] M. C. Miller *et al.*, PSR J0030 + 0451 mass and radius from NICER data and implications for the properties of neutron star matter, *Astrophys. J. Lett.* **887**, L24 (2019).
- [52] T. E. Riley *et al.*, A NICER view of PSR J0030 + 0451: Millisecond pulsar parameter estimation, *Astrophys. J. Lett.* **887**, L21 (2019).
- [53] M. C. Miller *et al.*, The radius of PSR J0740 + 6620 from NICER and XMM-Newton data, *Astrophys. J. Lett.* **918**, L28 (2021).
- [54] T. E. Riley *et al.*, A NICER view of the massive pulsar PSR J0740 + 6620 informed by radio timing and XMM-Newton spectroscopy, *Astrophys. J. Lett.* **918**, L27 (2021).
- [55] F. Douchin and P. Haensel, A unified equation of state of dense matter and neutron star structure, *Astron. Astrophys.* **380**, 151 (2001).
- [56] J. Margueron, R. Hoffmann Casali, and F. Gulminelli, Equation of state for dense nucleonic matter from meta-modeling. I. Foundational aspects, *Phys. Rev. C* **97**, 025805 (2018).
- [57] J. Margueron, R. Hoffmann Casali, and F. Gulminelli, Equation of state for dense nucleonic matter from meta-modeling. II. Predictions for neutron star properties, *Phys. Rev. C* **97**, 025806 (2018).
- [58] J. E. Lynn, I. Tews, J. Carlson, S. Gandolfi, A. Gezerlis, K. E. Schmidt, and A. Schwenk, Chiral three-nucleon interactions in light nuclei, neutron- α scattering, and neutron matter, *Phys. Rev. Lett.* **116**, 062501 (2016).
- [59] G. Baym, C. Pethick, and P. Sutherland, The ground state of matter at high densities: Equation of state and stellar models, *Astrophys. J.* **170**, 299 (1971).
- [60] E. Fonseca *et al.*, The NANOGrav nine-year data set: Mass and geometric measurements of binary millisecond pulsars, *Astrophys. J.* **832**, 167 (2016).
- [61] M. C. Miller, C. Chirenti, and F. K. Lamb, Constraining the equation of state of high-density cold matter using nuclear and astronomical measurements, [arXiv:1904.08907](https://arxiv.org/abs/1904.08907).
- [62] P. Landry, R. Essick, and K. Chatziioannou, Nonparametric constraints on neutron star matter with existing and upcoming gravitational wave and pulsar observations, *Phys. Rev. D* **101**, 123007 (2020).
- [63] P. T. H. Pang, L. Sivertsen, R. Somasundaram, T. Dietrich, S. Sen, I. Tews, M. Coughlin, and C. Van Den Broeck, Probing quarkyonic matter in neutron stars with the Bayesian nuclear-physics multi-messenger astrophysics framework, *Phys. Rev. C* **109**, 025807 (2024).
- [64] T. Gorda, K. Hebeler, A. Kurkela, A. Schwenk, and A. Vuorinen, Constraints on strong phase transitions in neutron stars, *Astrophys. J.* **955**, 100 (2023).
- [65] M. Ujevic, R. Somasundaram, T. Dietrich, J. Margueron, and I. Tews, What can we learn about the unstable equation-of-state branch from neutron-star mergers?, *Astrophys. J. Lett.* **962**, L3 (2024).
- [66] S. Tootle, C. Ecker, K. Topolski, T. Demircik, M. Järvinen, and L. Rezzolla, Quark formation and phenomenology in binary neutron-star mergers using V-QCD, *SciPost Phys.* **13**, 109 (2022).
- [67] G. D. Moore and T. Gorda, Bounding the QCD equation of state with the lattice, *J. High Energy Phys.* **12** (2023) 133.
- [68] Y. Fujimoto and S. Reddy, Bounds on the equation of state from QCD inequalities and lattice QCD, *Phys. Rev. D* **109**, 014020 (2024).
- [69] O. Komoltsev, T. Gorda, and A. Kurkela, Marginalized QCD likelihood function, Zenodo, [10.5281/zenodo.10592568](https://zenodo.org/record/10592568) (2024).
- [70] C. Drischler, R. J. Furnstahl, J. A. Melendez, and D. R. Phillips, How well do we know the neutron-matter equation of state at the densities inside neutron stars? A Bayesian approach with correlated uncertainties, *Phys. Rev. Lett.* **125**, 202702 (2020).
- [71] C. D. Capano, I. Tews, S. M. Brown, B. Margalit, S. De, S. Kumar, D. A. Brown, B. Krishnan, and S. Reddy, Stringent constraints on neutron-star radii from multimessenger observations and nuclear theory, *Nat. Astron.* **4**, 625 (2020).
- [72] M. Al-Mamun, A. W. Steiner, J. Nättilä, J. Lange, R. O’Shaughnessy, I. Tews, S. Gandolfi, C. Heinke, and S. Han, Combining electromagnetic and gravitational-wave constraints on neutron-star masses and radii, *Phys. Rev. Lett.* **126**, 061101 (2021).