



CHORUS

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

Toward a deeper understanding of how experiments constrain the underlying physics of heavy-ion collisions

Evan Sangaline and Scott Pratt

Phys. Rev. C **93**, 024908 — Published 10 February 2016

DOI: [10.1103/PhysRevC.93.024908](https://doi.org/10.1103/PhysRevC.93.024908)

Towards a Deeper Understanding of How Experiments Constrain the Underlying Physics of Heavy-Ion Collisions

Evan Sangaline and Scott Pratt

*Department of Physics and Astronomy and National Superconducting Cyclotron Laboratory
Michigan State University, East Lansing, MI 48824 USA*

(Dated: December 27, 2015)

Recent work has provided the means to rigorously determine properties of super-hadronic matter from experimental data through the application of broad scale modeling of high-energy nuclear collisions within a Bayesian framework. These studies have provided unprecedented statistical inferences about the physics underlying nuclear collisions by virtue of simultaneously considering a wide range of model parameters and experimental observables. Notably, this approach has been used to constrain both the QCD equation of state and the shear viscosity above the quark-hadron transition. Although the inferences themselves have a clear meaning, the complex nature of the relationships between model parameters and observables have remained relatively obscure. We present here a novel extension of the standard Bayesian Markov Chain Monte Carlo approach that allows for the quantitative determination of how inferences of model parameters are driven by experimental measurements and their uncertainties. This technique is then applied in the context of heavy ion collisions in order to explore previous results in greater depth. The resulting relationships are useful for identifying model weaknesses, prioritizing future experimental measurements, and most importantly: developing an intuition for the role that different observables play in constraining our understanding of the underlying physics.

I. INTRODUCTION

Relativistic nuclear collisions provide a unique opportunity to create and study matter at temperatures and energy densities well above the boundary where hadronic degrees of freedom become irrelevant and are replaced with quark and gluon degrees of freedom. Whereas the energy density inside a hadron might be on the order of $100 \text{ MeV}/\text{fm}^3$, high-energy nuclear collisions produce mesoscopic regions where the average energy density can exceed $10 \text{ GeV}/\text{fm}^3$. Unfortunately, they are far from ideal environments for extracting the properties of super-hadronic matter. Only the asymptotic momenta of outgoing particles are experimentally accessible, so three-dimensional models of the dynamics are essential for interpreting results and forming rigorous conclusions about either the nature and properties of the matter or about the evolution itself. One serious limitation is that the exact nature of the initial deposition and dispersion of energy, momentum, charge, and baryon number in heavy-ion collisions is currently only partially understood. Nearly any experimental measurement will be sensitive to these initial conditions as well as to combined contributions from the proceeding stages of rapid expansion and cooling through partially equilibrated QGP and hadronic phases. This entangles contributions from the various components of the underlying physics and obscures the interpretation of many experimental results.

The development of models that can reliably mimic the various complexities of heavy-ion collisions has been a key focus of the field and a necessary step towards condensing experimental results into quantitative statements about the underlying physics. Unfortunately, determining physical quantities like the shear and bulk viscosities of QGP is non-trivial even with high quality models in place. Experimental data can be largely described using these models but the conclusions drawn depend strongly on the validity of the initial state model and other model assumptions. This problem is only confounded by the computational complexity of typical heavy-ion collision models which, in many cases, makes a full exploration of model parameter space prohibitively expensive and, in turn, encourages simplifying assumptions.

The situation of dealing with very complicated models that map high dimensional parameter spaces to rich heterogeneous observational data sets is a common one. It is faced in disparate fields of science such as biology, econometrics, and cosmology. The common solution to this problem is to apply a Bayesian inference approach where a posterior distribution over model parameters can be mapped out based on their consistency with a set of observables using a Markov Chain Monte Carlo (MCMC) or a related approach [1, 2]. This allows for quantitative conclusions about the underlying parameters to be made given the validity of the model and prior assumptions.

A primary difficulty that has historically limited the applicability of a Bayesian approach in heavy-ion physics has been the large number of model evaluations necessary to accurately map out a high dimensional posterior distribution. This, coupled with the already computationally expensive nature of the simulations, makes a direct application of an MCMC approach essentially impossible with currently available hardware. Recent progress in methodology has overcome this difficulty through the use of Gaussian process emulators which can be trained to accurately reproduce the results of actual simulated models. The reduced computational load has opened the door to quantitative analyses of model parameters which has allowed for systematic constraints of the nuclear equation of state and of the QGP shear viscosity [3, 4] and details of the initial state [5].

Stating the single point in the multidimensional parameter space that maximizes the likelihood disregards the uncertainty with which the number is stated. In contrast, an MCMC exploration of parameter space provides not only the uncertainties for each parameter, but the covariances with other parameters along with the complete shape of the posterior likelihood. The methods of [3] indeed provide the complete likelihood distribution, and are now extracting the field's first rigorous quantitative conclusions from heavy-ion collisions at the highest energies.

Although the Bayesian approach has proved fruitful in nuclear physics, progress in understanding the complexities of generated posterior distributions and how they are driven by the experimental data has been slow. It is relatively straightforward to understand how parameters drive observables through varying them and seeing what changes. Going in the opposite direction, seeing how the measurements drive the inferences about parameters, is more challenging for several reasons. For example, changing a certain parameter might readily alter a given observable, but there may exist some combination of other parameters that can compensate without affecting the remaining observables. Even an observable which does not seem to directly depend on a given parameter, might help constrain other parameters, which then helps constrain the first parameter. In deciding which observables to either measure or to better measure, one would like to know how the width of the posterior likelihood distribution is affected by reducing the uncertainty of a given observable. These relations are critical to gaining insight into understanding not only the degree to which model parameters are being constrained, but how and why. Here, we present a new set of techniques for addressing all of these questions and then apply the techniques to the 14-parameter analysis first presented in [4]. This analysis represents the field's first quantitative evaluation of much of the field's heuristic understanding about which observables are truly responsible for addressing the community's most basic questions about the bulk properties of nuclear matter and about the evolution of a high-energy heavy-ion collision.

In the next two sections we review the model, data and techniques used to generate our previous results in [3, 4],

presenting a wider range of results from the projections of the MCMC procedure than were presented in [4]. In Sec. IV we present techniques for determining linear relations between observables and parameters and between uncertainties in observables and widths of the posterior distribution in parameter space. We apply these techniques to the heavy-ion analysis described in Sec.s II and III, and focus on the ramifications for extracting the equation of state and viscosity. Results are summarized and an outlook is presented in the final section.

II. MODEL AND DATA OVERVIEW

The details of the model are expounded in more detail in [3, 4], but we will briefly review the basics and focus on describing the 14 parameters varied in this analysis along with the set of observables. Our model consists of the generation of an initial state which is then fed into a 2-dimensional hydrodynamics simulation followed by transition to a microscopic simulation, known as a hadronic cascade, at a transition temperature of $T_0 = 165$ MeV. An assumed symmetry based on an invariance to boosts along the beam axis makes it possible to approximate three-dimensional treatment with a two-dimensional model that discretizes the two transverse coordinates.

Ten model parameters are used to vary the initial state at a time $\tau_0 = 0.8$ fm/c, two are used to determine the equation of state in the hydrodynamic stage, and the last two are used to determine the shear viscosity and its temperature dependence in the hydrodynamic stage. The shear viscosity is parameterized as

$$\frac{\eta}{s} = \left(\frac{\eta}{s}\right)_0 + \eta' \ln \frac{T}{T_0} \quad (1)$$

where $(\eta/s)_0$ is the viscosity at T_0 and η' describes the temperature dependence. These two parameters will be featured in this paper due to their relatively intuitive meaning and relationship with common observables. Two parameters also encapsulate the equation of state. The speed of sound, c_s , as a function of the energy density, ϵ , is described as

$$c_s^2(\epsilon) = c_s^2(\epsilon_h) + \left(\frac{1}{3} - c_s^2(\epsilon_h)\right) \frac{X_0 x + x^2}{X_0 x + x^2 + X'^2}, \quad (2)$$

$$X_0 = X' \Gamma c_s(\epsilon_h) \sqrt{12}, \quad x \equiv \ln(\epsilon/\epsilon_h),$$

where ϵ_h and $c_s(\epsilon_h)$ are the energy density and speed of sound of a hadron gas at the transition temperature, $T_0 = 165$ MeV. These quantities are calculated by considering a gas of non-interacting hadrons using the masses and spins of particles from the Particle Data Group [6]. All resonances with masses below 2 GeV/ c^2 were included. With this prescription, the equation of state is continuous at T_0 . The parameter Γ describe the behavior of c_s just above T_0 and the parameter X' provides a scale at which the speed of sound approaches $1/3$. Increasing X' lowers the speed of sound mainly at high energy density or temperature, and increasing Γ increases the speed of sound, mainly just above T_0 . This is illustrated with 5 equations of state in Fig. 1.

Five of the ten initial state parameters apply only to the description of RHIC (Relativistic Heavy Ion Collider at Brookhaven National Laboratory) data. Only data from Au+Au collisions at $100A$ GeV + $100A$ GeV is considered. The other five described the initial state for Pb+Pb collisions at $1.38A$ TeV + $1.38A$ TeV from the LHC (Large Hadron Collider at CERN). The initial transverse energy density profile that instantiates the hydrodynamics has the form [3]

$$\epsilon(x, y) = f_{\text{wn}} \epsilon_{\text{wn}}(x, y) + (1 - f_{\text{wn}}) \epsilon_{\text{sat}}(x, y), \quad (3)$$

where the parameter f_{wn} describes the weighting between two other parameterized forms, ϵ_{wn} which is the wounded nucleon form [7], and ϵ_{sat} , which describes a form more in line with some ideas of saturation, similarly to [8].

$$\epsilon_{\text{wn}}(x, y) = Z_\epsilon \frac{(dE_\perp/dy)_{pp} \sigma_{\text{nn}}}{2\sigma_{\text{sat}}} T_A(x, y) (1 - \exp(-T_B(x, y)\sigma_{\text{sat}})), \quad (4)$$

$$+ Z_\epsilon \frac{(dE/dy)_{pp} \sigma_{\text{nn}}}{2\sigma_{\text{sat}}} T_B(x, y) (1 - \exp(-T_A(x, y)\sigma_{\text{sat}})),$$

$$\epsilon_{\text{sat}}(x, y) = Z_\epsilon \frac{(dE/dy)_{pp} \sigma_{\text{nn}}}{\sigma_{\text{sat}}} T_{\text{min}}(x, y) (1 - \exp(-T_{\text{max}}(x, y)\sigma_{\text{sat}})), \quad (5)$$

$$T_{\text{min}} \equiv \frac{2T_A T_B}{T_A + T_B}, \quad T_{\text{max}} \equiv (T_A + T_B)/2.$$

Here T_A and T_B are the areal densities, the projections of the baryon density of a Au or Pb nucleus onto the transverse plane, and have dimensions of number per area. In the wounded nucleon model, each nucleon that comes within the




FIG. 1. (Color online) Equations of state, speed of sound vs. energy density, are shown for five sets of parameters. The solid black line is for a default case where both equation-of-state parameters X' and Γ from Eq. (2) are chosen from the center of the prior distribution given in Table I. The red dashed lines describe the cases where X' is unchanged from the default, but Γ is at the minimum/maximum of the prior distribution. This parameter most strongly affects the speed of sound just above $\epsilon_h \approx 0.5$ GeV/fm³. The case where Γ is set at the default values and X' is varied is represented by the blue dotted lines. The parameter X' varies the scale at which the speed of sound approaches $1/3$ asymptotically. Stiffer equations of state are attained by either raising Γ or lowering X' .

nucleon-nucleon cross section, $\sigma_{nn} = 42$ mb at RHIC energies and 73 mb at LHC energies, of one of the other nucleons, known as a participant, contributes to the energy density. If the parameter σ_{sat} is set equal to the nucleon-nucleon cross section then each nucleon can only contribute once and additional collisions do not increase the energy density. Relaxing $\sigma_{\text{sat}} < \sigma_{nn}$ allows the particle to contribute multiple times. As σ_{sat} approaches zero the form gives binary scaling and the resulting energy density is proportional to $T_A T_B$. In the saturated form, the energy density is principally determined by the smaller of the two areal densities. Thus, if one nucleon overlaps 5 other nucleons, the energy density will be only slightly less than if it overlapped 10.

The parameter Z_ϵ represents the energy per unit rapidity per nucleon collision in a dilute reaction relative to the measurement of a pp collision, and both forms satisfy the constraint that for a diffuse overlap of nucleons, i.e. $T_A, T_B \ll \sigma_{pp}$ the energy density scales as binary collisions, $\epsilon(x, y) \rightarrow \sigma_{pp} (dE/dy)_{pp} T_A T_B$ with $Z_\epsilon \approx 1$. However, because this is the energy density at $\tau_0 = 0.8$ fm/c, and not the final energy, and because the energy density includes longitudinal motion, it differs from the measured transverse energy, and Z_ϵ was allowed to vary over a small range near unity. For our purposes $(dE/dy)_{pp}$ was set to 2.69 and 6.0 GeV for the two beam energies, though this number is somewhat arbitrary because it is multiplied by Z_ϵ . Finally, the fifth parameter describes the initial transverse flow, which is assumed to have the form

$$\frac{T_{0i}}{T_{00}} = F_0 \frac{-\partial_i T_{00}}{2T_{00}} \tau, \quad (6)$$

where $T_{\alpha\beta}$ is the stress-energy tensor. For a traceless stress energy tensor, which would be expected for non-interacting gluon fields or collision-less massless particles or for conformal hydrodynamics, and also assuming boost invariance, F_0 would go to unity [9]. The final parameter adjusted the initial anisotropy of the stress-energy tensor,

$$T_{xx} = T_{yy} = P(1 + \tau'_{xx}), \quad T_{zz} = P(1 - 2\tau'_{xx}). \quad (7)$$

This form is traceless $(1/3) \sum_{i=1,2,3} T_{ii} = P$ and describes the initial shear, which in the Israel Stewart form of hydrodynamics is a dynamical variable which relaxes toward the Navier Stokes value. Parameters are summarized in Table I.

The model is used to produce a set of 30 observables that correspond to experimental measurements performed at RHIC and the LHC. These encompass distilled information describing spectra, elliptic flow, and femtoscopic correlations in central and mid-central collisions. These observables were chosen because they all have been shown to characterize features of thermalized bulk matter. Fifteen of the observables correspond to Au+Au collisions at RHIC while the other fifteen correspond to Pb+Pb collisions at the LHC. Observables were taken from two centrality classes, the top 0 – 5% centrality cut and the set of collisions in the 20 – 30% centrality cut. It was felt that additional data with centrality between these values would be redundant because of the smooth behavior of the observations with

Parameter	Min, Max (RHIC/LHC values)	Description
Γ	-0.9, 2.0	Equation of State
X'	0.5, 5.0	Equation of State
$(\eta/s)_0$	0.02, 0.5	Viscosity at T_0
η'	0.0, 3.0	Temperature dependence of viscosity
Z_ϵ	0.8 / 0.7, 1.25 / 1.4	Energy normalization ratio
$\sigma_{\text{sat}}(\text{mb})$	22 / 38, 44 / 76	Saturation cross section
f_{wn}	0.0 / 0.0, 1.0 / 1.0	Weight of wounded-nucleon parameterization
F_0	0.2 / 0.2, 1.0 / 1.0	Initial transverse flow
τ'_{xx}	0.0 / 0.0, 1.0 / 1.0	Initial anisotropy of stress-energy tensor

TABLE I. Fourteen model parameters

centrality. Further, it would be difficult to assign uncertainties for observables from more peripheral collisions due to the difficulty in justifying hydrodynamic treatments for systems whose size is not larger than a thermal wavelength, and where a significant fraction of the transverse profile is in the corona, which does not fully thermalize.

Three classes of measurements were considered. The first and most basic is spectra and yields. These were distilled to four numbers for each centrality. Three were the mean transverse momenta, p_t , for pions, kaons and protons. The average was taken over a finite range of p_t , limited at the low end by experimental constraints and cut off at the high end to minimize the effects of jets which are non-thermal features outside of this analysis. The fourth measurement was the multiplicity of pions within the same p_t range. Because chemical equilibrium was assumed, rather than parameterized, and because baryon annihilation was not included, the yields of kaons and protons were not included. In [3] it was shown that model runs that had the same $\langle p_t \rangle$ had the same spectral shapes, so no resolving power was lost by considering only one number to characterize the spectra. Those spectra with the same $\langle p_t \rangle$ as the data also provided good descriptions of the experimental spectral shapes [3, 4].

The second class of observables is comprised of femtoscopic radii extracted from two-pion correlations at small relative momentum. Extracting these radii has long been a staple of the field [10]. The gaussian radii are functions of momentum, and describe the size and shape of the outgoing phase space cloud of the given momentum. The three radii, R_{out} , R_{side} and R_{long} , describe the transverse dimension parallel to the momentum, the transverse size perpendicular to the momentum and the size along the beam axis respectively. For this analysis, the radii were averaged over the various bins in transverse momentum, so that three observables encapsulated the data for each centrality class.

The final observable was v_2 , which characterizes the anisotropic transverse flow which is driven by the elliptical shape of the initial transverse profile. The anisotropy, $v_2 \equiv \langle \cos 2\phi \rangle$ in its simplest definition, was only considered in the 20-30% centrality bin. This bin was chosen because the model used smooth profiles, generated from the average aerial profiles for a given impact parameter, and neglected the lumpy conditions which one would expect from the finite number of scatterers in a Au or Pb nucleus. Even though this bin is probably the least affected by lumps, the experimental value was reduced by a factor of 0.91 to more fairly compare to the model [4]. The next major analysis planned by this group will incorporate fluctuating initial conditions, and in addition to more accurately calculating v_2 , could also address the higher components $v_3, v_4 \dots$, which are mainly driven by fluctuations. Because v_2 rises nearly linearly with transverse momentum, the value of v_2 was averaged over p_t , with the values weighted by p_t , so that bins with higher p_t were relatively more important. The linear weighting was motivated by performing a principal component analysis of the binned values to find the combination that best captured the variability in the model runs. Table II summarizes the 30 observables in this analysis.

III. MARKOV CHAIN MONTE CARLO PROCEDURE AND RESULTS

The standard approach to determining the likely regions of parameters, \vec{x} , from comparing model values, $\vec{y}_M(\vec{x})$, to experimental values, \vec{y}_{exp} , is through MCMC, which provides a sampling of parameters \vec{x} that are chosen weighted proportional to the likelihood. For this study the likelihood is chosen to have a simple Gaussian form,

$$\mathcal{L}(\vec{x}) \sim \exp \left\{ - \sum_a \frac{(y_{M,a}(\vec{x}) - y_a^{(\text{exp})})^2}{2\sigma_a^2} \right\}. \quad (8)$$

The uncertainties incorporate both experimental uncertainties and the shortcomings of the model. For instance, if the equation of state and viscosity were perfectly described, and if the initial state was parameterized most correctly,

observable	exp. value	p_t weighting	centrality	collaboration
$v_{2,\pi^+\pi^-}$	8.14%	ave. over 11 p_t bins from 160 MeV/c to 1 GeV/c	20-30%	STAR [11]
R_{out}	5.28 fm	ave. over 4 p_t bins from 150-500 MeV/c	0-5%	STAR [12]
R_{side}	4.81 fm	ave. over 4 p_t bins from 150-500 MeV/c	0-5%	STAR [12]
R_{long}	5.47 fm	ave. over 4 p_t bins from 150-500 MeV/c	0-5%	STAR [12]
R_{out}	4.27 fm	ave. over 4 p_t bins from 150-500 MeV/c	20-30%	STAR [12]
R_{side}	3.99 fm	ave. over 4 p_t bins from 150-500 MeV/c	20-30%	STAR [12]
R_{long}	4.53 fm	ave. over 4 p_t bins from 150-500 MeV/c	20-30%	STAR [12]
$\langle p_t \rangle_{\pi^+\pi^-}$	494.4 MeV	0.2 GeV/c < p_t < 1.2 GeV/c	0-5%	PHENIX [13]
$\langle p_t \rangle_{K^+K^-}$	796 MeV	0.4 GeV/c < p_t < 1.6 GeV/c	0-5%	PHENIX [13]
$\langle p_t \rangle_{p\bar{p}}$	1.135 GeV	0.6 GeV/c < p_t < 2.0 GeV/c	0-5%	PHENIX [13]
$\langle p_t \rangle_{\pi^+\pi^-}$	487.5 MeV	0.2 GeV/c < p_t < 1.2 GeV/c	20-30%	PHENIX [13]
$\langle p_t \rangle_{K^+K^-}$	792 MeV	0.4 GeV/c < p_t < 1.6 GeV/c	20-30%	PHENIX [13]
$\langle p_t \rangle_{p\bar{p}}$	1.111 GeV	0.6 GeV/c < p_t < 2.0 GeV/c	20-30%	PHENIX [13]
$\pi^+\pi^-$ yield	422	0.2 GeV/c < p_t < 1.2 GeV/c	0-5%	PHENIX [13]
$\pi^+\pi^-$ yield	188.7	0.2 GeV/c < p_t < 1.2 GeV/c	20-30%	PHENIX [13]
$v_{2,\pi^+\pi^-}$	9.56%	ave. over 11 p_t bins from 0.15 to 1.2 GeV/c	20-30%	ALICE [14]
R_{out}	5.46 fm	ave. over 7 p_t bins from 200 to 900 MeV/c	0-5%	ALICE [15]
R_{side}	5.32 fm	ave. over 7 p_t bins from 200 to 900 MeV/c	0-5%	ALICE [15]
R_{long}	5.72 fm	ave. over 7 p_t bins from 200 to 900 MeV/c	0-5%	ALICE [15]
R_{out}	4.07 fm	ave. over 7 p_t bins from 200 to 900 MeV/c	20-30%	ALICE [15]
R_{side}	4.12 fm	ave. over 7 p_t bins from 200 to 900 MeV/c	20-30%	ALICE [15]
R_{long}	4.41 fm	ave. over 7 p_t bins from 200 to 900 MeV/c	20-30%	ALICE [12]
$\langle p_t \rangle_{\pi^+\pi^-}$	459.1 MeV	0.1 GeV/c < p_t < 1.2 GeV/c	0-5%	ALICE [16]
$\langle p_t \rangle_{K^+K^-}$	775 MeV	0.2 GeV/c < p_t < 1.6 GeV/c	0-5%	ALICE [16]
$\langle p_t \rangle_{p\bar{p}}$	1.137 GeV	0.2 GeV/c < p_t < 2.0 GeV/c	0-5%	ALICE [16]
$\langle p_t \rangle_{\pi^+\pi^-}$	455 MeV	0.1 GeV/c < p_t < 1.2 GeV/c	20-30%	ALICE [16]
$\langle p_t \rangle_{K^+K^-}$	758 MeV	0.2 GeV/c < p_t < 1.6 GeV/c	20-30%	ALICE [16]
$\langle p_t \rangle_{p\bar{p}}$	1.110 MeV	0.2 GeV/c < p_t < 2.0 GeV/c	20-30%	ALICE [16]
$\pi^+\pi^-$ yield	1258	0.2 GeV/c < p_t < 1.2 GeV/c	0-5%	ALICE [16]
$\pi^+\pi^-$ yield	523	0.2 GeV/c < p_t < 1.2 GeV/c	20-30%	ALICE [16]

TABLE II. Observables used to compare models to data from RHIC (STAR and PHENIX collaborations) and from ALICE (at the LHC). To account for non-flow correlations, the value of v_2 was reduced by 9% from the experimental values to account for the approximation of smooth initial conditions.

the model uncertainty describes how accurately one would expect to reproduce the final-state observables given the missing physics in the model and the uncertainties in the measurements. In this study, which uses the same model output as in [4], the total uncertainties σ_a were all set to 6% of the experimental value $y_a^{(\text{exp})}$. A more accurate determination of the uncertainty would require detailed model studies to estimate the impact of missing physics, and a more detailed understanding of experimental uncertainties. It would not be surprising to find that a few of these observables are understood with slightly better uncertainty or that some are somewhat more uncertain. Such an improvement should involve input from both the experimental and modeling communities. In [3] the analysis was repeated with uncertainties changed by a factor of two and the widths of the resulting posterior distributions only increased by $\approx 20\%$.

Due to the infeasibility of performing millions of full model runs, the MCMC procedure implemented an emulator in the place of the full model [2, 3]. The emulator uses an interpolation algorithm to determine the observables from 1200 full-model runs. The first 1000 parameter sets for the full-model runs uniformly covered the 14-dimensional space according to latin hyper-cube sampling. The last 200 parameter sets were chosen to be consistent with the likelihood as calculated from the earlier runs and to provide better coverage in the likely region. To improve efficiency, the emulator calculates principal components rather than each observable. In this way, those linear combinations of observables that stay constant throughout the model runs can be neglected. Thus, rather than evaluating all 30 observables, only 14 principal components were analyzed for this study. Even that number could be reduced because the last several components only varied by a few percent of one value of the experimental uncertainty.




FIG. 2. (Color online) One- and two-dimensional posterior likelihood projections from the MCMC procedure. The one-dimensional projections (along the diagonal) illustrate the degree to which an individual parameter is constrained by comparing to data, whereas the off-diagonal elements illustrate how some linear combinations of two parameters might be either poorly or well constrained. The colored boundaries delineate one-, two- and three-sigma regions, where n -sigma refers to likelihoods of $e^{-n^2/2}$ of the maximum likelihood. The ranges for each parameter are listed in Table I.

The MCMC procedure was performed with several million random steps according to the Metropolis algorithm. This provides a sampling of the posterior distribution displayed in Fig. 2. In addition to projections onto one dimension of the parameter space, the off-diagonal plots show two-dimensional projections. When the elliptic shapes of the two-dimensional projections lie at an angle, it shows that a specific linear combination of parameters may be well constrained, whereas the orthogonal combination may be poorly constrained. For example, the region of high likelihood for the projection onto the plane of two equation of state parameters, one can see that if one increases




FIG. 3. (Color online) Posterior likelihood distributions for the shear viscosity parameters generated using RHIC and LHC data separately (a & b) and together (c). We can see that the viscosity at T_0 is well constrained by RHIC data alone but that there is very little constraint of its temperature dependence without the LHC data. This is consistent with the fact that LHC data probe higher temperatures. Combining the datasets clearly constrains the data better than a single data set. Panels (d-e) show corresponding information for the two equation of state parameters. In this case it is clear that LHC data provide the bulk of the resolving power. Again, this is expected because the equation of state is fixed at $T_0 = 165$ MeV, and LHC collisions probe regimes further into the region where the equation of state differs across the parameter space. One can also see that the high degree of covariance in both figures (c) and (f), especially for the covariance between the two equation-of-state parameters in (f). This shows that even though neither parameter is particularly well constrained, there is a linear combination that is well constrained, while the orthogonal linear combination is poorly constrained. Limits for each parameter are provided in Table I.

Γ and X' by similar percentages that the likelihood changes little. Similarly, for the two viscosity parameters one finds that one can find a good match with $(\eta/s)_0 \approx 0.2$ with a modest temperature dependence, η' . One can also match with lower values of $(\eta/s)_0$ if the viscosity then rises with temperature. Another instance where the posterior is off-diagonal is in the projection of $(\eta/s)_0$ and f_{wn} . The resulting covariance corroborates the arguments that were put forward in [8].

IV. SENSITIVITY STUDIES

The principal goal of this paper is to present various methods for understanding the role certain observables, or sets of observables, play in constraining the posterior likelihood of given parameters, or sets of parameters. The most straight-forward method to determine the sensitivity is to perform the analysis both with and without a given observable, or set of observables. An example of this is illustrated in Fig. 3. The two-dimensional posterior likelihood projection for the two viscosity parameters, $(\eta/s)_0$ and η' , are redisplayed along with projections where either RHIC data or LHC data are ignored. This makes it clear that the LHC data is especially important for constraining the temperature dependence of the viscosity, η' .

However, the method of repeating the statistical analysis without a given observable can be onerous. If one wishes to study the sensitivity to $N_Y = 30$ observables, one would repeat the MCMC procedure N_Y times. Here, we present methods for finding three simple measures of the sensitivity using the output from a single MCMC trace. That information would be of the form of a list of values, $\alpha = 1 \cdots N_{\text{MCMC}}$, where N_{MCMC} might be of the order of one million. For each point one would have the parameters $x_{\alpha,i}$ and the observables $y_{\alpha,a}$.

The first measure would simply describe how a given observable, y_a , would change due to a change in a parameter x_i , while keeping all other parameters constant. Because $y_{M,a}^{(\bar{x})}$ is not purely linear, one needs to choose over what region the derivative is evaluated. Two choices of interest might be the prior or the posterior. Here, we use single brackets, $\langle \cdots \rangle$, to denote an average of the prior, and double brackets, $\langle\langle \cdots \rangle\rangle$, to denote averages taken over the posterior. For averages over the prior, one can use the lists from the points at which full-model runs were performed for the purpose of training the emulator. If calculations were performed randomly throughout the prior, one could consider the covariance $\langle \delta y_a \delta x_i \rangle$, where $\delta y_a = y_a - \langle y_a \rangle$ and $\delta x_i = x_i - \langle x_i \rangle$. One can calculate the partial derivative of y_a with respect to any parameter by a simple matrix inversion,

$$\langle \delta y_a \delta x_i \rangle = \sum_j \left\langle \frac{\partial y_a}{\partial x_j} \right\rangle \langle \delta x_j \delta x_i \rangle \quad (9)$$

$$\left\langle \frac{\partial y_a}{\partial x_j} \right\rangle = \langle \delta y_a \delta x_i \rangle \langle \delta x \delta x \rangle_{ij}^{-1}.$$

Here, the brackets around $\partial y / \partial x$ note that this is the best slope for this region of parameters space, which in this case is the prior. This expression matches the usual least-squares expression for the slope of a line in multidimensions. Equation 9 can be easily altered to address how the slope looks when focused on the posterior region by making the change $\langle \cdots \rangle \rightarrow \langle\langle \cdots \rangle\rangle$.

A. Response of extracted parameter values to changes in experimental measurement

Equation 9 shows how a specific model value, y_a , changes when a specific parameter, x_i , is varied. However, describing how a small change in an experimental measurement, $y_a^{(\text{exp})}$, affects the average value of a parameter, $\langle\langle x_i \rangle\rangle$, in the average posterior value is a different question. This involves understanding $\partial\langle\langle x_i \rangle\rangle/\partial y_a^{(\text{exp})}$. Here we consider a general expression for any posterior-averaged function of \vec{x} , $\langle\langle f(\vec{x}) \rangle\rangle$, which for this specific consideration will be $f(\vec{x}) = x_i$.

$$\begin{aligned} \frac{\partial}{\partial y_a^{(\text{exp})}} \langle\langle f(\vec{x}) \rangle\rangle &= \frac{\partial}{\partial y_a^{(\text{exp})}} \frac{\int d\vec{x} f(\vec{x}) \mathcal{L}(\vec{x})}{\int d\vec{x} \mathcal{L}(\vec{x})} \\ &= \langle\langle f(\vec{x}) \frac{1}{\mathcal{L}} \frac{\partial}{\partial y_a^{(\text{exp})}} \mathcal{L} \rangle\rangle - \langle\langle f(\vec{x}) \rangle\rangle \langle\langle \frac{1}{\mathcal{L}} \frac{\partial}{\partial y_a^{(\text{exp})}} \mathcal{L} \rangle\rangle. \end{aligned} \quad (10)$$

Both terms on the right-hand-side of Eq. (10) can be calculated from the trace. By setting $f = \delta x_i = (x_i - \bar{x}_i)$, with $\bar{x}_i \equiv \langle\langle x_i \rangle\rangle$, the second term vanishes in this case, and if one has an expression for the likelihood, one need only average $\delta x_i (1/\mathcal{L}) \partial \mathcal{L}$ over the sampling of points from the MCMC trace to calculate the required result. One can invoke the emulator to calculate $y(\vec{x})$, and thereby determine both \mathcal{L} and $\partial \mathcal{L}$.

For Gaussian likelihoods Eq. (10) becomes especially simple. In that case

$$\begin{aligned} \mathcal{L} &= \exp \left\{ \frac{1}{2} (y_a - y_a^{(\text{exp})}) \Sigma_{ab}^{-1} (y_b - y_b^{(\text{exp})}) \right\} \\ \frac{1}{\mathcal{L}} \frac{\partial}{\partial y_a^{(\text{exp})}} \mathcal{L} &= \Sigma_{ab}^{-1} (y_a - y_a^{(\text{exp})}), \\ \frac{\partial}{\partial y_a^{(\text{exp})}} \bar{x}_i &= \Sigma_{ab}^{-1} \langle\langle \delta x_i \delta y_b \rangle\rangle, \\ &= \Sigma_{ab}^{-1} \left\langle \left\langle \frac{\partial y_a}{\partial x_j} \right\rangle \right\rangle A_{ji}, \\ A_{ij} &\equiv \langle\langle \delta x_i \delta x_j \rangle\rangle. \end{aligned} \quad (11)$$

Here, δy can be relative to any point because $\langle\langle \delta x \rangle\rangle = 0$, and the derivative in the last line is the average slope for the posterior region. Eq. (11) is straightforward to calculate from the output of the MCMC. For the case where Σ is diagonal,

$$\left. \frac{\partial \bar{x}_i}{\partial y_a^{(\text{exp})}} \right|_{y_b^{(\text{exp})}} = \frac{1}{\sigma_a^2} \langle\langle \delta y_a \delta x_i \rangle\rangle \quad (12)$$

To better quantify how a specific observable constrains a given parameter, one may multiple the expression for $\partial \bar{x}/\partial y^{(\text{exp})}$ in Eq. (12) by the measure of how much y_a changes throughout the prior, $\langle\langle \delta y_a^2 \rangle\rangle^{1/2}$, because model values of an observable must change within the range of the prior if that observable is to provide resolving power. Figure 4 displays this quantity for all observables. One can also see whether the change is positive or negative, which is not obvious and may have an opposite sign compared to $\partial y/\partial x$.

As examples of the sensitivity analysis one can look at the elliptic flow measured at RHIC and at the LHC. In Fig. 5 one can see that changing the measurement of v_2 at RHIC strongly affects the extracted value of the viscosity at T_0 , $(\eta/s)_0$, but has little effect on the extracted temperature dependence, η' . In contrast, the measurement of v_2 at the LHC more strongly affects η' , while playing a more minor role in determining $(\eta/s)_0$. Given the fact that the LHC explores higher energy densities, this was expected. One can also see that the constraints on the viscosity were driven by measurements of v_2 and the multiplicities, whereas constraints on the two equation-of-state parameters were driven by a wider variety of measurements.

B. Relation between experimental uncertainties and widths of posterior parameter distributions

We now consider how the uncertainty of a specified observable affects the width of the posterior parameter distribution, i.e., calculate the derivative

$$\frac{\partial}{\partial \Sigma_{ab}} A_{ij} = \langle\langle \delta x_i \delta x_j \frac{1}{\mathcal{L}} \frac{\partial \mathcal{L}}{\partial \Sigma_{ab}} \rangle\rangle - \langle\langle \delta x_i \delta x_j \rangle\rangle \langle\langle \frac{1}{\mathcal{L}} \frac{\partial \mathcal{L}}{\partial \Sigma_{ab}} \rangle\rangle, \quad (13)$$




FIG. 4. (Color online) Model responses of an observable with respect to a given parameter, $\langle\langle \partial y_a / \partial x_i \rangle\rangle / \sigma_a$.

where the last step followed the steps from the previous subsection used to derive Eq. (10). This provides a similar measure of how the parameters are resolved by specific observables, and to that end provides very similar usefulness as the calculations of $\partial \bar{x}_i / \partial y_a^{(\text{exp})}$ displayed in the right-hand side of Fig. ???. Whereas the previous measure described how the mean of the posterior distribution for a given parameter, x_i , changes when one changes a given observable, the quantity in Eq. (13) provides an understanding of how changing the uncertainty for a given observable, σ_a , affects the width of the posterior distribution for a given parameter. One expects the two measures to provide similar insight.

Equation (13) is fairly straight-forward to calculate, but can be a bit cumbersome. To find a simpler, though approximate, expression we assume that the prior distribution of parameters is Gaussian even though the results shown here all assumed hard cutoffs, and that $y(\vec{x})$ is linear, so that the posterior distribution is Gaussian. We also use Gaussian likelihoods for comparing to experimental values. When comparing calculations with and without these




FIG. 5. (Color online) The change of the inferred value of a parameter with respect to changes in a measurement, $\partial\langle\langle x_i \rangle\rangle/\partial y_a^{(\text{exp})}$, are scaled by the spread of model values throughout the prior, $\langle\delta y_a^2\rangle^{1/2}$. Larger absolute values point to measurements which play important roles in constraining that parameter. The signs of the response are not always equal for the corresponding derivatives in this plot and in Fig. 4.

approximations, results changed little.

With the approximations, the overall likelihood is then of a Gaussian form,

$$\mathcal{L} \sim \exp \left\{ -|\delta x|^2/2R^2 - \left\langle \left\langle \frac{\partial y_a}{\partial x_i} \right\rangle \right\rangle \Sigma_{ab}^{-1} \left\langle \left\langle \frac{\partial y_b}{\partial x_j} \right\rangle \right\rangle \frac{\delta x_i \delta x_j}{2} \right\}, \quad (14)$$

$$\delta x_i = x_i - \bar{x}_i.$$

We have assumed that the parameters x_i have all been scaled to have the same prior width R . These parameters thus

become dimensionless. Further, in the neighborhood of the maximum likelihood we assume that y behaves linearly. The width of the posterior, A_{ij} , can be taken from the Gaussian form,

$$A_{ij} = \frac{\int d^N x \exp\{-A_{ij}^{-1} \delta x_i \delta x_j / 2\} \delta x_i \delta x_j}{\int d^N x \exp\{-A_{ij}^{-1} \delta x_i \delta x_j / 2\}},$$

$$A_{ij}^{-1} \equiv \delta_{ij} / R^2 + \left\langle \left\langle \frac{\partial y_a}{\partial x_i} \right\rangle \right\rangle \Sigma_{ab}^{-1} \left\langle \left\langle \frac{\partial y_b}{\partial x_j} \right\rangle \right\rangle. \quad (15)$$

Our stated goal is to find an expression describing how A_{ij} responds to changes in Σ_{ab} ,

$$\begin{aligned} \frac{\partial A}{\partial \Sigma_{ab}^{-1}} &= \frac{\partial}{\partial \Sigma_{ab}^{-1}} (A A^{-1} A) \\ &= 2 \frac{\partial}{\partial \Sigma_{ab}^{-1}} A + A \left(\frac{\partial}{\partial \Sigma_{ab}^{-1}} A^{-1} \right) A \\ &= -A \left(\frac{\partial}{\partial \Sigma_{ab}^{-1}} A^{-1} \right) A \\ \frac{\partial}{\partial \Sigma_{ab}^{-1}} A_{i\ell} &= -A_{ij} \left\langle \left\langle \frac{\partial y_a}{\partial x_j} \right\rangle \right\rangle \left\langle \left\langle \frac{\partial y_b}{\partial x_k} \right\rangle \right\rangle A_{k\ell}. \end{aligned} \quad (16)$$

$$(17)$$

One can now insert the expression for $\partial y / \partial x$ given in Eq. (9) into Eq.s 15 and 17 to find the needed result. When working with the hard cutoffs for priors rather than Gaussians, we scale all the parameters to have the same prior width, $\langle x^2 \rangle = R^2$, in addition to being centered at zero, $\langle x \rangle = 0$. If one uses the prior distribution to calculate $\partial y / \partial x$, then the response in Eq. (17) need not ever access the experimental value, and one does not need to perform the MCMC.

If one wishes to use the posterior distribution to calculate the derivatives, the form for the response simplifies

$$\frac{\partial}{\partial \Sigma_{ab}^{-1}} A_{ij} = -\langle \langle \delta x_i \delta y_a \rangle \rangle \langle \langle \delta y_b \delta x_j \rangle \rangle, \quad (18)$$

which can also be transformed into an expression in terms of derivatives with respect to Σ ,

$$\begin{aligned} \frac{\partial}{\partial \Sigma_{ab}} A_{ij} &= \frac{\partial A_{ij}}{\partial \Sigma_{cd}^{-1}} \frac{\partial \Sigma_{cd}^{-1}}{\partial \Sigma_{ab}} \\ \frac{\partial \Sigma_{cd}^{-1}}{\partial \Sigma_{ab}} &= \frac{\partial}{\partial \Sigma_{ab}} \left(\Sigma_{ce}^{-1} \Sigma_{ef} \Sigma_{fd}^{-1} \right) \\ &= -\Sigma_{ca}^{-1} \Sigma_{bd}^{-1}, \\ \frac{\partial}{\partial \Sigma_{ab}} A_{ij} &= \Sigma_{ac}^{-1} \langle \langle \delta y_c \delta x_i \rangle \rangle \langle \langle \delta y_d \delta x_j \rangle \rangle \Sigma_{db}^{-1}. \end{aligned} \quad (19)$$

In this analysis the uncertainty matrix is diagonal, $\Sigma_{ab} = \sigma_a^2 \delta_{ab}$, and Eq. (19) can be used to calculate the resolving power for determining a single parameter x_i due to a single parameter y_a .

$$\begin{aligned} R_{i;a} &\equiv \frac{\sigma_a}{\langle (\delta x_i)^2 \rangle} \frac{\partial}{\partial \sigma_a} A_{ii}, \\ &= \frac{\sigma_a}{\langle (\delta x_i)^2 \rangle} \Sigma_{ac}^{-1} \langle \langle \delta y_c \delta x_i \rangle \rangle \langle \langle \delta y_d \delta x_j \rangle \rangle \Sigma_{da}^{-1}. \end{aligned} \quad (20)$$

Results for $R_{i;a}$ are shown in Fig. 6.

Even when a specific parameter is not well constrained, a combination of that parameter with another might still be well constrained. The two equation-of-state parameters are a good example, and the two viscosity parameters also have a strong off-diagonal component to their likelihood projections. These two instances are shown in panels (c) and (f) in Fig. 3. If one then wants to understand the contribution of a given observable in constraining the two-dimensional projection, one needs to define a quantity which effectively measures the likely area of the two-dimensional projection onto the parameters x_i and x_j . Here we use the determinant of the two-by-two matrix constructed from the ij components of A_{ij} .

$$D_{ij} \equiv A_{ii} A_{jj} - A_{ij} A_{ji}, \quad (21)$$

$$R_{ij;a} \equiv \sigma_a \frac{\partial}{\partial \sigma_a} D_{ij} = \sigma_a \left\{ A_{ii} \frac{\partial}{\partial \sigma_a} A_{jj} + A_{jj} \frac{\partial}{\partial \sigma_a} A_{ii} - 2 A_{ij} \frac{\partial}{\partial \sigma_a} A_{ji} \right\}. \quad (22)$$




FIG. 6. (Color online) The resolving power of a particular observable in determining a specific parameter, $R_{i;a}$ defined in Eq. (20), is displayed for all observables and parameters.

D_{ij} would represent the product of the eigenvalues of the two-by-two matrix, or equivalently, the square of the area covered by the projection. The way in which D_{ij} changes with respect to a given uncertainty can then be calculated with the help of Eq. (19). These sensitivities are illustrated in Fig. 7. This clearly demonstrates that the extracted viscosity is strongly affected by the v_2 measurements, and that multiplicities are also important. Other observables are of secondary, but not negligible importance. For constraining the equation of state, femtoscopic radii seem to provide the most resolving power, but all other observables contribute significantly. This underscores the importance of simultaneously considering multiple classes of observables to constrain the parameter space.

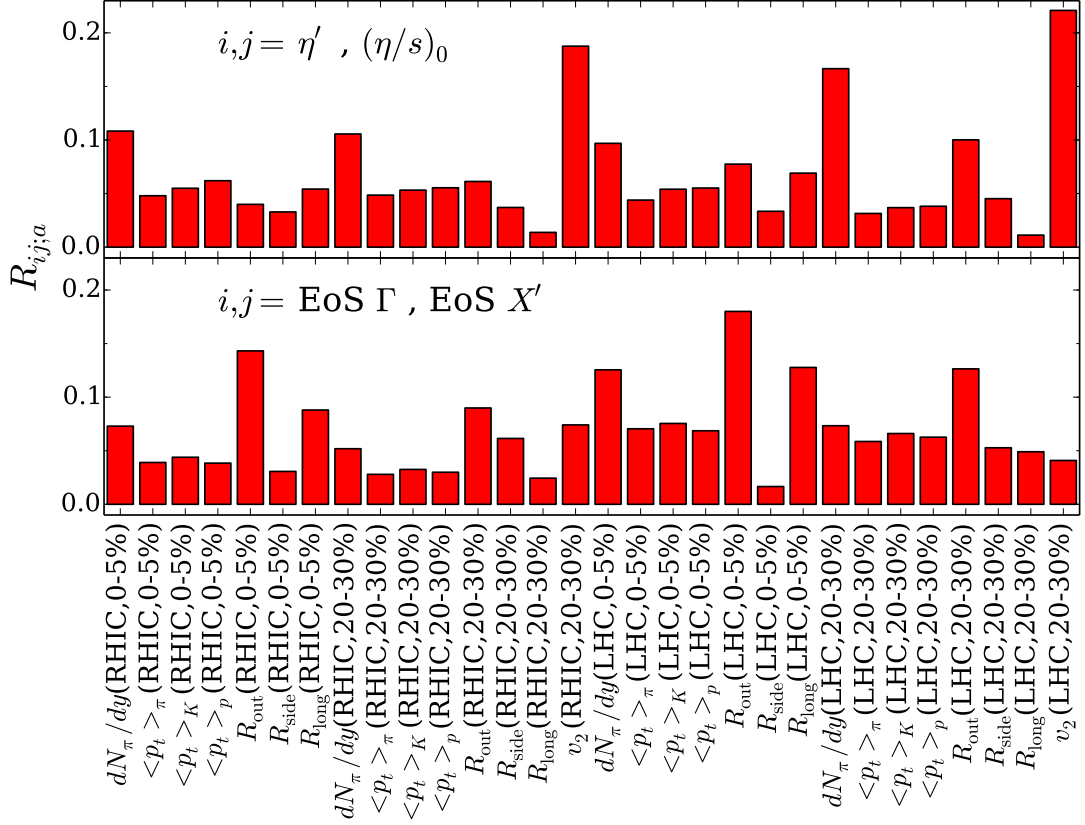


FIG. 7. (Color online) The resolving power for determining a pair of parameters, x_i, x_j , due to a specific observable y_a as defined by $R_{ij;a}$ in Eq. (22). The sensitivity for the two viscosity parameters is shown in the upper panel and that for the two equation-of-state parameters is displayed in the lower panels. Information about the viscosity is most strongly determined by measurements of v_2 and the multiplicities. The equation of state is most strongly constrained by femtosopic radii.

V. SUMMARY AND OUTLOOK

The results here are unprecedented for this field, and for the first time illustrate a systematic method for identifying the critical links between parameters and observables. The links identified by the procedure mainly reinforced the heuristically attained general knowledge of the field. For instance, it was indeed found that measurements of the elliptic anisotropy provide strong constraints of the viscosity. Further, the expectation that RHIC data would play a larger relative role in determining the viscosity near T_c , and that the LHC data would play a larger role in determining the temperature dependence was confirmed. The relatively large role of the multiplicities was not necessarily expected, nor was the fact that other observables provide non-negligible resolving power.

Whereas there was consensus within the field that v_2 would be the most important observable to determine the viscosity, the role of various observables in constraining the equation of state was very much in dispute. Figures 6, ?? and 7 all show that the femtosopic source sizes provide the most resolving power. This validates ideas based on low pressures leading to more elongated phase space distributions in the outward direction and that the longitudinal sizes would increase if transverse expansion was slower, [10, 17, 18]. Additionally, source sizes played a role in measuring the final entropy, which is also a measure of the equation of state [19]. The mean transverse momentum had long ago been pointed out as being sensitive to the temperature, and therefore the equations of state, [20], and even v_2 had been suggested as a means for extracting early pressure [21]. Therefore, it was not surprising to see the resolving power for the equation of state in Fig. 7 spread across all observables.

The techniques presented here could play a pivotal role in determining the direction of future experiments. Before embarking on an expensive experimental program to improve the measure of a specific observable, one could check to see how that improvement might indeed better determine parameters of greatest interest. For example, one could determine whether running an accelerator for a specific projectile target combination, or at a new beam energy, or with higher statistics would best provide insight into better determining the equation of state. In many cases it would behoove the community to pre-analyze a project with modern statistical techniques before investing the cost and

manpower for the effort.

ACKNOWLEDGMENTS

This work was supported by the National Science Foundation's Cyber-Enabled Discovery and Innovation Program through grant NSF-0941373 and by the Department of Energy Office of Science through grant number DE-FG02-03ER41259.

-
- [1] S. Habib, K. Heitmann, D. Higdon, C. Nakhleh and B. Williams, *Phys. Rev. D* **76**, 083503 (2007).
 - [2] F. A. Gomez, C. E. Coleman-Smith, B. W. O'Shea, J. Tumlinson and R. L. Wolpert, *Astrophys. J.* **760**, 112 (2012).
 - [3] J. Novak, K. Novak, S. Pratt, J. Vredevoogd, C. E. Coleman-Smith and R. L. Wolpert, *Phys. Rev. C* **89**, 034917 (2014).
 - [4] S. Pratt, E. Sangaline, P. Sorensen and H. Wang, *Phys. Rev. Lett.* **114**, 202301 (2015).
 - [5] J. E. Bernhard, P. W. Marcy, C. E. Coleman-Smith, S. Huzurbazar, R. L. Wolpert and S. A. Bass, *Phys. Rev. C* **91**, no. 5, 054910 (2015) [arXiv:1502.00339 [nucl-th]].
 - [6] Particle Data Group, *The Review of Particle Physics*, <http://pdg.lbl.gov> (2012).
 - [7] For an overview of wounded nucleon models see M. L. Miller, K. Reygers, S. J. Sanders and P. Steinberg, *Ann. Rev. Nucl. Part. Sci.* **57**, 205 (2007).
 - [8] H. -J. Drescher, A. Dumitru, C. Gombeaud and J. -Y. Ollitrault, *Phys. Rev. C* **76**, 024905 (2007).
 - [9] J. Vredevoogd and S. Pratt, *Phys. Rev. C* **79**, 044915 (2009).
 - [10] M. A. Lisa, S. Pratt, R. Soltz and U. Wiedemann, *Ann. Rev. Nucl. Part. Sci.* **55**, 357 (2005).
 - [11] J. Adams *et al.* [STAR Collaboration], *Phys. Rev. C* **72**, 014904 (2005).
 - [12] B. I. Abelev *et al.* [STAR Collaboration], *Phys. Rev. C* **80**, 024905 (2009).
 - [13] S. S. Adler *et al.* [PHENIX Collaboration], *Phys. Rev. C* **69**, 034909 (2004).
 - [14] B. B. Abelev *et al.* [ALICE Collaboration], *JHEP* **1506**, 190 (2015).
 - [15] K. Graczykowski [ALICE Collaboration], *EPJ Web Conf.* **71**, 00051 (2014).
 - [16] B. Abelev *et al.* [ALICE Collaboration], *Phys. Rev. C* **88**, 044910 (2013).
 - [17] S. Pratt, *Phys. Rev. D* **33**, 1314 (1986).
 - [18] S. Pratt, *Phys. Rev. Lett.* **102**, 232301 (2009).
 - [19] S. Pal and S. Pratt, *Phys. Lett. B* **578**, 310 (2004).
 - [20] L. Van Hove, *Phys. Lett. B* **118**, 138 (1982).
 - [21] H. Sorge, *Phys. Rev. Lett.* **78**, 2309 (1997).