

UC Berkeley

UC Berkeley Previously Published Works

Title

Multisystem Bayesian constraints on the transport coefficients of QCD matter

Permalink

<https://escholarship.org/uc/item/1nn8108g>

Journal

Physical Review C, 103(5)

ISSN

2469-9985

Authors

Everett, D

Ke, W

Paquet, J-F

et al.

Publication Date

2021-05-01

DOI

10.1103/physrevc.103.054904

Peer reviewed

Multi-system Bayesian constraints on the transport coefficients of QCD matter

D. Everett,¹ W. Ke,^{2,3} J.-F. Paquet,⁴ G. Vujanovic,⁵ S. A. Bass,⁴ L. Du,¹ C. Gale,⁶ M. Heffernan,⁶ U. Heinz,¹ D. Liyanage,¹ M. Luzum,⁷ A. Majumder,⁵ M. McNelis,¹ C. Shen,^{5,8} Y. Xu,⁴ A. Angerami,⁹ S. Cao,⁵ Y. Chen,^{10,11} J. Coleman,¹² L. Cunqueiro,^{13,14} T. Dai,⁴ R. Ehlers,^{13,14} H. Elfner,^{15,16,17} W. Fan,⁴ R. J. Fries,^{18,19} F. Garza,^{18,19} Y. He,²⁰ B. V. Jacak,^{2,3} P. M. Jacobs,^{2,3} S. Jeon,⁶ B. Kim,^{18,19} M. Kordell II,^{18,19} A. Kumar,⁵ S. Mak,¹² J. Mulligan,^{2,3} C. Nattrass,¹³ D. Oliinychenko,³ C. Park,⁶ J. H. Putschke,⁵ G. Roland,^{10,11} B. Schenke,²¹ L. Schwiebert,²² A. Silva,¹³ C. Sirimanna,⁵ R. A. Soltz,^{5,9} Y. Tachibana,⁵ X.-N. Wang,^{20,2,3} and R. L. Wolpert¹²

(The JETSCAPE Collaboration)

¹*Department of Physics, The Ohio State University, Columbus OH 43210.*

²*Department of Physics, University of California, Berkeley CA 94270.*

³*Nuclear Science Division, Lawrence Berkeley National Laboratory, Berkeley CA 94270.*

⁴*Department of Physics, Duke University, Durham NC 27708.*

⁵*Department of Physics and Astronomy, Wayne State University, Detroit MI 48201.*

⁶*Department of Physics, McGill University, Montréal QC H3A 2T8, Canada.*

⁷*Instituto de Física, Universidade de São Paulo, C.P. 66318, 05315-970 São Paulo, SP, Brazil.*

⁸*RIKEN BNL Research Center, Brookhaven National Laboratory, Upton NY 11973.*

⁹*Lawrence Livermore National Laboratory, Livermore CA 94550.*

¹⁰*Laboratory for Nuclear Science, Massachusetts Institute of Technology, Cambridge MA 02139.*

¹¹*Department of Physics, Massachusetts Institute of Technology, Cambridge MA 02139.*

¹²*Department of Statistical Science, Duke University, Durham NC 27708.*

¹³*Department of Physics and Astronomy, University of Tennessee, Knoxville TN 37996.*

¹⁴*Physics Division, Oak Ridge National Laboratory, Oak Ridge TN 37830.*

¹⁵*GSI Helmholtzzentrum für Schwerionenforschung, 64291 Darmstadt, Germany.*

¹⁶*Institute for Theoretical Physics, Goethe University, 60438 Frankfurt am Main, Germany.*

¹⁷*Frankfurt Institute for Advanced Studies, 60438 Frankfurt am Main, Germany.*

¹⁸*Cyclotron Institute, Texas A&M University, College Station TX 77843.*

¹⁹*Department of Physics and Astronomy, Texas A&M University, College Station TX 77843.*

²⁰*Key Laboratory of Quark and Lepton Physics (MOE) and Institute of Particle Physics, Central China Normal University, Wuhan 430079, China.*

²¹*Physics Department, Brookhaven National Laboratory, Upton NY 11973.*

²²*Department of Computer Science, Wayne State University, Detroit MI 48202.*

(Dated: November 10, 2020)

We study the properties of the strongly-coupled quark-gluon plasma with a multistage model of heavy ion collisions that combines the $T_{\text{R}}\text{ENTO}$ initial condition ansatz, free-streaming, viscous relativistic hydrodynamics, and a relativistic hadronic transport. A model-to-data comparison with Bayesian inference is performed, revisiting assumptions made in previous studies. The role of parameter priors is studied in light of their importance towards the interpretation of results. We emphasize the use of closure tests to perform extensive validation of the analysis workflow before comparison with observations. Our study combines measurements from the Large Hadron Collider and the Relativistic Heavy Ion Collider, achieving a good simultaneous description of a wide range of hadronic observables from both colliders. The selected experimental data provide reasonable constraints on the shear and the bulk viscosities of the quark-gluon plasma at $T \sim 150\text{--}250$ MeV, but their constraining power degrades at higher temperatures $T \gtrsim 250$ MeV. Furthermore, these viscosity constraints are found to depend significantly on how viscous corrections are handled in the transition from hydrodynamics to the hadronic transport. Several other model parameters, including the free-streaming time, show similar model sensitivity, while the initial condition parameters associated with the $T_{\text{R}}\text{ENTO}$ ansatz are quite robust against variations of the particlization prescription. We also report on the sensitivity of individual observables to the various model parameters. Finally, Bayesian model selection is used to quantitatively compare the agreement with measurements for different sets of model assumptions, including different particlization models and different choices for which parameters are allowed to vary between RHIC and LHC energies.

CONTENTS

		1. Energy deposition at $\tau = 0^+$: $T_{\text{R}}\text{ENTO}$	5
		2. Free-streaming	6
I. Introduction	2	B. Relativistic viscous hydrodynamics	7
		C. Particlization	8
II. Inference using Bayes' Theorem	3	1. Linear viscous corrections: Grad & Chapman-Enskog	9
III. Model Overview	5	2. Exponentiated viscous corrections: Pratt-Torrieri-McNelis and	
A. Initial stage model	5		

Pratt-Torrieri-Bernhard	10
D. Hadronic transport	11
IV. Specifying prior knowledge	11
V. Bayesian Parameter Estimation with a Statistical Emulator	13
A. Overview of Bayesian Parameter Estimation	13
B. Physical model emulator	14
C. Treatment of uncertainties	16
D. Sampling of the posterior	17
E. Maximizing the posterior	17
VI. Closure Tests	17
A. Validating Bayesian inference with closure tests	18
B. Guiding analyses with closure tests	18
VII. Bayesian parameter estimation using RHIC and LHC measurements	19
A. Constraints on η/s and ζ/s from Pb-Pb measurements at $\sqrt{s_{NN}} = 2.76$ TeV	19
B. Constraints on η/s and ζ/s from Au-Au measurements at $\sqrt{s_{NN}} = 0.2$ TeV	20
C. Viscosity estimation and model accuracy for combined RHIC & LHC data	21
VIII. Parameter estimation and systematic model uncertainties	22
A. Mapping hydrodynamic fields to hadronic momentum distributions	22
B. Transition to and from hydrodynamics: initial state and switching temperature	23
C. Second-order transport coefficients: shear relaxation time	26
IX. Sensitivity to prior knowledge and assumptions	27
X. Model Sensitivity	27
XI. Bayesian Model Selection	29
A. Overview and estimation of Bayes factors	30
1. Bayes factor definition and interpretation	30
2. Numerical methods for estimating the Bayes evidence	30
B. Comparing viscous correction models	30
C. Comparing hydrodynamic models	32
1. Temperature independent specific shear viscosity	32
2. Zero specific shear viscosity	32
D. Quantifying tension between LHC and RHIC	32
1. No common parameters between collision systems	32
2. Different transverse length scales in the initial conditions	33
XII. Predicting p_T -differential observables	33
XIII. Summary	34
XIV. Outlook	36

Acknowledgments	37
A. Full posterior of model parameters	37
B. Posterior for LHC and RHIC independently	37
C. Validation of principal component analysis	37
D. Experimental covariance matrix	38
E. Reducing experimental uncertainty	39
F. Bulk relaxation time	39
G. Comparison to previous studies	40
1. Physics models	41
2. Prior distributions	42
3. Experimental data	42
H. Multistage model validation	42
1. Validation of second-order viscous hydrodynamics implementation	42
a. Validation against cylindrically symmetric external solution	43
2. SMASH	43
3. Comparison of JETSCAPE with hic-eventgen	45
4. The σ meson	46
5. Sampling particles on mass-shell	47
6. QCD equations of state with different hadron resonance gases	47
References	48

I. INTRODUCTION

One of the primary goals of the heavy ion program pursued at the Relativistic Heavy Ion Collider (RHIC) and the Large Hadron Collider (LHC) is a quantitative understanding of the many-body properties of nuclear matter under extreme conditions as described by Quantum Chromodynamics (QCD). Lattice QCD calculations [1, 2] predict that hot and dense nuclear matter undergoes a cross-over transition near a temperature of 150 MeV at vanishing net baryon density, changing from hadronic degrees of freedom to a deconfined phase of nuclear matter — the quark-gluon plasma (QGP). While lattice QCD calculations nowadays provide very precise first-principles information on the equilibrium properties of QCD matter across this transition from hadrons to color-deconfined QGP [1, 2], the transport properties of QCD matter, such as its shear and bulk viscosities (which are of critical importance for its collective dynamical behavior in heavy-ion collisions), so far remain under the purview of phenomenology.¹ Calculating the viscosities of QCD from first principles remains challenging, in particular in the range of temperatures reached in heavy-ion

¹ See Ref. [3] for a recent overview.

collisions ($T \sim 150\text{--}500$ MeV) where non-perturbative effects are important.²

The distributions and correlations of hadrons produced in heavy-ion collisions provide a complementary approach to quantify these transport properties of a QCD medium in the laboratory. Multistage dynamical models of heavy-ion collisions based on relativistic viscous hydrodynamics have been successful in describing a wide range of soft ($p_T \lesssim 2$ GeV) hadronic observables from RHIC and LHC [12–15]. It was realized early on that hadronic observables measured in heavy-ion collisions are sensitive to the shear [16, 17] and bulk [18, 19] viscosities of QCD matter. Early works in constraining these transport coefficients with hydrodynamic simulations of heavy ion collisions generally focused on the shear viscosity (*c.f.* the reviews [12, 13]) which was approximated as having a constant ratio to the entropy density η/s (i.e. a constant *specific* shear viscosity). Contemporary efforts attempt to constrain the temperature and chemical potential dependence of the specific shear and bulk viscosities, namely $(\eta/s)(T, \mu_B)$ and $(\zeta/s)(T, \mu_B)$ [20, 21]. Large-scale model-to-data comparisons are necessary to achieve this goal, given the considerable challenge of constraining the full functional form of $(\eta/s)(T, \mu_B)$ and $(\zeta/s)(T, \mu_B)$ in the relevant temperature and chemical potential ranges.

The primary goal of this work is to explore phenomenological constraints on the temperature dependence of both the specific shear and bulk viscosities of the nuclear matter produced in relativistic heavy ion collisions at high energy, where net baryon density is negligible at midrapidity. We use a modern multistage model to simulate heavy ion collisions and perform a Bayesian parameter estimation on all the associated model parameters, including shear and bulk viscosities. The current analysis focuses on soft hadron measurements in the central rapidity region, as the available data are precise and they are arguably the observables whose theoretical uncertainties are best under control.

This work builds upon recent Bayesian analyses of soft hadronic observables [22–26], in particular Ref. [27]. One difference of the current analysis with these previous works is our use of a more flexible parametrization of the model; our more general functional forms for $(\eta/s)(T)$ and $(\zeta/s)(T)$ and the wider range of viscosity values allowed (“the priors”) in the model-to-data comparison plays an important role in this new analysis. The role of the emulator is also discussed in greater depth, in particular the importance of closure tests as a validation of the Bayesian approach. For the first time, we quantify the uncertainties of $(\eta/s)(T)$ and $(\zeta/s)(T)$ stemming from known theoretical uncertainties in modeling viscous corrections to particlization; in the accompanying Letter [28] we use Bayesian model averaging to arrive at improved estimates by combining the constraints from different particlization models. The effects of the shear relaxation time on our constraints on $(\eta/s)(T)$ and $(\zeta/s)(T)$ are also quantified. We further ex-

plore constraining $(\eta/s)(T)$ and $(\zeta/s)(T)$ by combining experimental measurements from both RHIC and the LHC, and explore the subtleties of including measurements from different experiments.

Whereas one of the features of our work is the exploration of theoretical uncertainties associated with the viscous corrections to particlization, and attempts to minimize the latter by calibrating the model with p_T -integrated flow observables [29], an orthogonal approach is pursued in the recent works [30, 31] which restrict their analysis to a single particlization model and instead explore additional gains in precision by including additional p_T -differential flow observables.

Our work is performed within the greater scope of the JETSCAPE Collaboration [32]. The physics objectives for the Collaboration are (i) to obtain a complete and reliable space-time description of the quark-gluon plasma by performing a state-of-the-art calibration of its key properties using the large body of available experimental data, and (ii) to use this calibrated dynamical evolution model for performing precision studies of penetrating and hard probes of the evolving medium. Our focus here is on the first objective.

Unless noted otherwise, all equations in this manuscript use natural units ($\hbar = c = k_B = 1$). We use the “mostly-minus” metric signature: $g^{\mu\nu} = \text{diag}(1, -1, -1, -1)$ in Cartesian coordinates. In general, position and momentum four-vectors are denoted with capital letters, and their components with Greek indices, e.g. the space-time position four-vector X has components X^μ ; three-vectors are denoted with boldface, and their components with Latin indices, e.g. the spatial position three-vector \mathbf{X} has components X^i . The present work focuses on midrapidity observables from high-energy heavy-ion collisions where longitudinal boost-invariance is a good approximation. Such systems are most efficiently implemented using Milne coordinates $X^\mu = (\tau, \mathbf{x}_\perp, \eta_s)$, with $\tau = \sqrt{t^2 - z^2}$ and $\eta_s = \frac{1}{2} \ln[(t+z)/(t-z)]$.

II. INFERENCE USING BAYES’ THEOREM

During the evolution of heavy ion collisions, the systems probe a wide range of many-body regimes of Quantum Chromodynamics. Because of the microscopic system size and its ultra-fast dynamics, this many-body physics must be inferred from the only information experimentally accessible in heavy ion collisions: the energy-momentum spectra and correlations of final-state particles that hit the detectors. These particles include stable (under the strong interaction) hadrons such as pions, kaons, protons, unstable hadronic resonances, and also electroweak bosons.

Given the complicated and non-linear correlations between parameters of the many-body dynamics and the observables, it is rare that a given physical effect — for example the quark-gluon plasma having a certain temperature-dependent viscosity — can be associated with a single observable constructed from final state particles. In general, quite different physical scenarios for the various evolution stages of the collisions may lead to quantitatively similar predictions for any single final state observable. Disentangling different physical scenarios

² See for example Ref. [4] and references therein for recent efforts at evaluating the QGP shear viscosity on the lattice, and Refs. [5–11] for other approaches to calculating these viscosities.

requires simultaneous analysis of large sets of complementary observables, each having been measured with finite uncertainties. This is a formidable challenge, known as the “inverse problem” of complex models. It makes the field of heavy ion collisions a natural candidate for the application of advanced statistical methods to study many-body QCD.

From an abstract point of view, systematically performing statistical inference with a complex model requires a formalism or set of axioms that can form the basis for “plausible reasoning”. The rules of Bayesian probability offer a natural formalism to systematically tackle such inverse problems. A Bayesian definition of the conditional probability of some proposition A given known information B , denoted $\mathcal{P}(A|B)$, is a quantification of our degree of belief about A , or its “plausibility”. A simple but powerful identity, which follows from the product rule of probability, is Bayes’ theorem:

$$\mathcal{P}(A|B) = \frac{\mathcal{P}(B|A)\mathcal{P}(A)}{\mathcal{P}(B)}. \quad (1)$$

Here the probability $\mathcal{P}(A)$ represents our prior knowledge or belief about the proposition A , $\mathcal{P}(B|A)$ is the likelihood for B to be true if the proposition A holds, and the normalization $\mathcal{P}(B) = \int dA \mathcal{P}(B|A)\mathcal{P}(A)$ is called the overall “Bayes evidence” for the information B to be true. The left hand side of Eq. (1) is known as the posterior probability distribution (“posterior” in short) for the proposition A given the information B . Theorem (1) allows us to invert the order of conditioning when we want to perform plausible reasoning or inference about some proposition A with knowledge B in hand. A more thorough introduction to Bayesian inference can be found in Ref. [33].

Broadly, we could quantify many different propositions. For example, one may want to quantify the likely values of transport coefficients such as shear and bulk viscosity, discussed in more detail later, given measured values for the soft ($p_T \lesssim 2$ GeV) hadronic observables such as multiplicities or mean transverse momenta. Another example of immediate interest is quantifying the likely values of transport coefficients controlling the energy-momentum exchange of hard partons and the quark-gluon plasma, given hard ($p_T \gtrsim 8$ GeV) observables such as the nuclear modification factor of inclusive hadron production. Any proposition regarding the physics of heavy ion collisions can in principle be quantified using Bayesian inference.

When performing inference for heavy-ion collisions, we only have the experimentally measured observables, complemented by first-principles physics considerations, to guide us towards what could plausibly be the dynamics of the collision. To make quantitative statements about the quark-gluon plasma requires a dynamical model for the entire heavy ion collision evolution. Broadly, the model is the relation between quantities of interest (such as the viscosities) and the observables: a “map” from the model parameters to the observables.

Many widely-used models for heavy ion collisions include a stage of evolution that is described by hydrodynamics, where dynamical properties are specified by a set of transport coefficients, such as the shear and bulk viscosities. Thus, a major goal of the phenomenology of heavy ion physics has been to in-

fer the viscosities of the QGP given hydrodynamic models for its expansion. However, only the embedding of these macroscopic hydrodynamic models within a more sophisticated class of “hybrid models” [34–42], which add modules for describing on a more microscopic level the early pre-hydrodynamic and late hadronic rescattering and freeze-out stages, has enabled a fully quantitative modeling of heavy-ion collisions. The quantitative predictive power and precision of these modern dynamical approaches have now opened the door for serious efforts to tackle the “inverse problem”: instead of providing a range of model predictions based on a limited scan of the model parameter space using subjective selection criteria, the community is now beginning to exploit the full set of available experimental observations to provide quantitative estimates for the model parameters, with quantified uncertainties.

The method of inferring the likely model parameters given the observed data with use of Bayes’ theorem is called Bayesian parameter estimation. This is the natural tool for estimating the likely values of the viscosities of the QGP, given a hydrodynamic model. This will occupy the largest portion of this work, beginning with a description of the hybrid model in Section III and a description of parameter estimation in Sections IV through VIII. In Section X, we take a brief aside to examine a measure of our model’s sensitivity to changes in the parameters. This helps us understand which of the observables carry information about parameters, thereby also allowing us to interpret the effects of different observables on the posterior.

Evidently, the likely values of any parameters depend not only on the observed data, but also on the model at hand. Some parameters that have meaning in a given model may not be relevant for a different model. Therefore, in general, it does not suffice to take a single model that works reasonably well, perform parameter estimation, and conclude from the estimations that the system is quantitatively understood. It is likely that there exists a “superior” model. In this context “superiority” should not only be measured by how well a model can fit a set of data. With sufficient complexity, any model can begin to overfit the data from a given experiment. Such overfitting leads to a model having a less universal applicability: when the model is applied to make a prediction in a previously unexplored region, it can perform very poorly. The best model to explain a given set of data should balance “simplicity” with accuracy in reproducing the experimental data.

A statistical tool for judging the relative merit of two models is the Bayes factor (the ratio of their Bayesian evidences as defined in Eq. (1) which weighs both their ability to fit measurements and their simplicity. Choosing between competing models using Bayes factors is called Bayesian model selection. We will use the Bayes factor to compare various choices of hydrodynamic hybrid models for heavy ion collisions in Section XI. It can happen that, after performing model selection for competing models that share certain common parameters, none of the models studied is strongly favored over all others by the available experimental data. In this situation, one might want to estimate the posterior for the shared parameters by averaging over the considered models [43]. This procedure of Bayesian Model Averaging is explored in [28], where it is used to obtain estimates for the shear and bulk viscosity that include

the known theoretical model uncertainty regarding particle production, i.e. the process of converting hydrodynamic output into momentum spectra for finally observed particles.

III. MODEL OVERVIEW

Measurements of soft ($p_T \lesssim 2$ GeV) particles from heavy ion collisions can be described well by multistage dynamical simulations whose core is relativistic viscous hydrodynamics. Features of many-body QCD enter hydrodynamics through medium properties such as the equation of state and transport coefficients (e.g. the shear and bulk viscosities). The collision dynamics preceding the applicability of hydrodynamics are described separately with a “pre-hydrodynamic stage” model. Following the hydrodynamic evolution, the last stage of the collision proceeds microscopically via hadronic transport.

The multistage model used in this work combines the following ingredients:

- initial energy deposition from the colliding nuclei, given by the T_RENTO ansatz [44, 45], followed by free-streaming expansion [46–48];
- relativistic viscous hydrodynamic evolution [49–52] employing a lattice QCD based equation of state [2, 53, 54] and flexible temperature dependent parametrizations of the first-order transport coefficients;
- conversion of the nuclear fluid into particles employing the Cooper-Frye approach, where several different models are used for mapping the fluid’s energy-momentum tensor to hadronic momentum distributions [55, 56];
- hadronic rescatterings in the hadronic transport model SMASH [57, 58] until kinetic freeze-out.

In the following subsections we present details on each of these ingredients.

A. Initial stage model

First-principles descriptions of the pre-hydrodynamic stage of a heavy ion collision remain challenging. However, many microscopic details of this early stage of the collisions are thought to be irrelevant [59] for the initialization of hydrodynamics at a time τ_{is} of order 0.1–1 fm/c (which, in kinetic theory language, requires only the first and second momentum moments of the microscopic distribution function). We therefore parametrize the initial conditions of hydrodynamics at τ_{is} , assuming that the system got to this time from $\tau = 0^+$ by free-streaming. Evidence from other microscopic theories of the early dynamics in heavy-ion collisions [60–64] suggests that (at least for systems that are sufficiently weakly coupled to admit a kinetic theory description of their microscopic dynamics [65]) coarse-grained (i.e. long wavelength) features of the energy-momentum tensor of a conformal system follow a relatively simple evolution similar to the free-streaming approach employed here. The energy deposition at $\tau = 0^+$ is

parametrized with the T_RENTO ansatz. This approach should be judged by the flexibility of the combination of T_RENTO with free-streaming, rather than by each component individually. Both components are now described in more detail.

1. Energy deposition at $\tau = 0^+$: T_RENTO

We use Woods-Saxon profiles to describe the distribution of nucleons in each colliding heavy nucleus.³ In the plane transverse to the collision axis the energy⁴ deposition immediately following the impact between the two nuclei is parametrized with the T_RENTO ansatz. This model parametrizes the transverse energy deposition via a reduced thickness function T_R ,

$$\bar{\epsilon}(\mathbf{x}_\perp) \equiv \frac{dE}{d\eta d^2\mathbf{x}_\perp} = NT_R(\mathbf{x}_\perp; p), \quad (2)$$

$$T_R(\mathbf{x}_\perp; p) = \left(\frac{T_A^p(\mathbf{x}_\perp) + T_B^p(\mathbf{x}_\perp)}{2} \right)^{1/p}, \quad (3)$$

where N is a normalization parameter estimated by comparison with data, and T_A and T_B represent the participant nucleon areal densities of the two nuclei.⁵ In Eq. (2) we assume longitudinal boost invariance for the collision system, and $\bar{\epsilon}(\mathbf{x}_\perp) \equiv \lim_{\tau \rightarrow 0^+} \tau \epsilon(\tau, \mathbf{x}_\perp, \eta=0)$. The continuous parameter p defines a family of mappings from T_A and T_B to the energy deposition. Specific values of p are known to reproduce certain features of other initial condition models [44]. When $p \rightarrow 0$, $T_R \rightarrow \sqrt{T_A T_B}$ and the model shares similar relations between eccentricities (ϵ_2, ϵ_3) and centrality [44] as those predicted by the IP-Glasma initial condition model [68]. This $\sqrt{T_A T_B}$ -scaling is also expected from imposing the conservation of longitudinal momentum with a flux-tube profile ansatz for energy density along space-time rapidity [69]. For $p = 1$, the energy deposition is equivalent to the wounded nucleon Glauber model [70]. Nucleon positions are sampled from the Woods-Saxon profiles of the heavy nuclei; a parameter d_{min} controls the minimum distance between any pair of nucleons to mimic the short range repulsion of the nucleon potential.⁶ Nuclear collisions are generated by performing binary-nucleon inelastic collisions in a Monte-Carlo procedure.⁷ The collision probability for two nucleons separated by an impact parameter

³ The Woods-Saxon parameters used for Au are radius $R = 6.38$ fm and surface thickness $a = 0.535$ fm while those for Pb are $R = 6.62$ fm and $a = 0.546$ fm.

⁴ While initial studies used T_RENTO to parametrize entropy density [26, 44, 66] we follow later works [27, 67] which use an energy density parametrization before free-streaming.

⁵ $T_{A/B}$ should not be confused with the similar nuclear thickness function that often appears in the optical Glauber model.

⁶ By default T_RENTO sets a fixed value for d_{min} but we treat the corresponding volume d_{min}^3 as variable when performing our parameter estimation. In particular, we will assign a uniform prior for d_{min}^3 .

⁷ We define minimum bias events in T_RENTO as events with at least a single binary collision. Events that produce no particles because they have no switching hypersurface into the hydrodynamic phase are included with zero multiplicity in the centrality selection.

b is

$$\frac{dP}{d^2b} = 1 - \exp \left[-\sigma_{gg} \int d^2\mathbf{x}_\perp \rho \left(\mathbf{x}_\perp + \frac{\mathbf{b}}{2} \right) \rho \left(\mathbf{x}_\perp - \frac{\mathbf{b}}{2} \right) \right]. \quad (4)$$

$\sigma_{gg}(\sqrt{s})$ is an effective nucleon opacity fixed by the proton-proton inelastic cross-section at a given beam energy,

$$\sigma_{pp}^{\text{in}}(\sqrt{s}) = \int d^2b \frac{dP}{d^2b}(\mathbf{b}, \sigma_{gg}), \quad (5)$$

and ρ is the longitudinally-integrated nucleon density per unit area in the transverse plane,

$$\rho(\mathbf{x}_\perp) = \int_{-\infty}^{\infty} \frac{dz}{(2\pi w^2)^{3/2}} \exp \left(-\frac{\mathbf{x}_\perp^2 + z^2}{2w^2} \right), \quad (6)$$

i.e. the nucleon thickness function. Here, we have modeled the density distribution of a nucleon as a three-dimensional Gaussian function with a width parameter w . Finally, each participant nucleon contributes to the participant density function,

$$T_A(\mathbf{x}_\perp) = \sum_{i \in A} \gamma_i \rho(\mathbf{x}_\perp - \mathbf{x}_{i,\perp}), \quad (7)$$

and similarly for nucleus B , where $\mathbf{x}_{i,\perp}$ is the position of participant nucleon in the transverse plane, and each nucleon's contribution is individually modulated by a stochastic factor γ_i sampled from a Γ -distribution with unit mean and standard deviation parameter σ_k .⁸ This additional source of fluctuations is introduced to model the large multiplicity fluctuations observed in minimum-bias proton-proton collisions.

2. Free-streaming

In this study, a non-trivial initial energy-momentum tensor is generated via a free-streaming model. It takes the initial energy density profile by T_RENTo and dynamically generates non-zero components for the entire energy-momentum tensor. This section provides details pertaining to how this is achieved.

The energy density (2) generated with T_RENTo at $\tau_0 = 0$ is assumed to represent massless degrees of freedom with a locally isotropic momentum distribution $f(X; P)$ centered at zero transverse momentum whose effective temperature (or mean P_T) varies with position X . The initial phase-space distribution is evolved by free-streaming, i.e. it is assumed to solve the collisionless Boltzmann equation

$$P^\mu \partial_\mu f(X; P) = 0, \quad (8)$$

which has the solution

$$f(t, \mathbf{x}; P) = f(t_0, \mathbf{x} - \mathbf{v}(t - t_0); P). \quad (9)$$

For massless particles $P^0 = |\mathbf{P}|$ and $\hat{p}^\mu \equiv P^\mu / P^0 = (1, \mathbf{v})$ where $|\mathbf{v}| = c$, i.e. all particles move with the speed of light. It

is convenient to define a moment $F(X; \Omega_p)$ of the distribution function by

$$F(X; \Omega_p) \equiv \frac{g}{(2\pi)^3} \int P_0^3 dP_0 f(X; P) \quad (10)$$

where g is a degeneracy factor. Then the stress tensor at any time $t > t_0$ is given by

$$\begin{aligned} T^{\mu\nu}(t, \mathbf{x}) &= \int d\Omega_p \hat{p}^\mu \hat{p}^\nu F(t, \mathbf{x}; \Omega_p) \\ &= \int d\Omega_p \hat{p}^\mu \hat{p}^\nu F(t_0, \mathbf{x} - \mathbf{v}(t - t_0); \Omega_p), \end{aligned} \quad (11)$$

where Ω_p is the solid angle in momentum space. The second equality was obtained by inserting the free-streaming solution (9).

Since we assume that the initial distribution function is locally isotropic in momentum space, the moment $F(t_0)$ can be related to the initial energy density by normalization:

$$T^{tt}(t_0, \mathbf{x}) = \mathcal{N} F(t_0, \mathbf{x}). \quad (12)$$

In three spatial dimensions $\mathcal{N} = 4\pi$. Here we assume longitudinal boost invariance, rendering the momentum distribution essentially two-dimensional, and therefore use $\mathcal{N} = 2\pi$ in the mapping (12). A more detailed description of the boost invariant case can be found in [46–48].

Free-streaming is stopped at a longitudinal proper time τ_{fs} , and the free-streamed energy-momentum tensor (11) is matched to that of viscous hydrodynamics through Landau matching conditions. The energy density ϵ in the local rest frame (LRF) and the flow velocity u^μ are the eigenvalue and timelike eigenvector of $T^{\mu\nu}$:

$$u_\mu T^\mu_\nu = \epsilon u_\nu. \quad (13)$$

The shear stress tensor is also matched exactly and given by the traceless and transverse projection of the stress tensor:

$$\pi^{\mu\nu} = \Delta^{\mu\nu}_{\alpha\beta} T^{\alpha\beta}. \quad (14)$$

where the transverse-traceless projector is defined by

$$\Delta^{\mu\nu}_{\alpha\beta} \equiv \frac{1}{2} (\Delta^\mu_\alpha \Delta^\nu_\beta + \Delta^\nu_\alpha \Delta^\mu_\beta) - \frac{1}{3} \Delta^{\mu\nu} \Delta_{\alpha\beta} \quad (15)$$

and

$$\Delta^{\mu\nu} \equiv g^{\mu\nu} - u^\mu u^\nu \quad (16)$$

projects onto the spatial directions in the local rest frame. The LRF is defined by $u^\mu_{\text{LRF}} = (1, \mathbf{0})$. Because the free-streaming dynamics continuously drive the system out of equilibrium, the size of the initial shear stress tensor grows with the free-streaming time τ_{fs} . We have checked in our numerical calculations that the second-order viscous hydrodynamics can reliably evolve the free-streaming $\pi^{\mu\nu}$ tensor to near local equilibrium for $\tau_{\text{fs}} \lesssim 1.5 \text{ fm}/c$.

The combination of thermal and bulk pressure, $p + \Pi$, is given by

$$p(\epsilon) + \Pi = \frac{\epsilon - T^\mu_\mu}{3}. \quad (17)$$

⁸ By default T_RENTo sets a fixed value for $k = 1/\sigma_k^2$, but we treat σ_k as variable when performing our parameter estimation.

Since we assume massless degrees of freedom (i.e. conformal symmetry) during the free-streaming stage, the energy-momentum tensor (11) is traceless, $T^\mu{}_\mu = 0$. The QGP fluid described by viscous hydrodynamics, on the other hand, is characterized by an equation of state that breaks conformal symmetry by interactions and non-zero quark masses: $p_{\text{QCD}}(\epsilon) < \epsilon/3$. Matching of the free-streamed energy-momentum tensor to the hydrodynamic one thus entails a non-zero, positive initial value for the fluid's bulk viscous pressure at τ_{fs} :

$$\Pi = \frac{\epsilon}{3} - p_{\text{QCD}}(\epsilon) \geq 0. \quad (18)$$

Note that for an expanding system a negative bulk viscous pressure is expected in the Navier-Stokes limit; the hydrodynamic evolution therefore quickly erases the positive initial value (18) and switches its sign. Persistence effects from the initially positive bulk viscous pressure depend on the value of the bulk relaxation time, which is studied in Appendix F. Different pre-hydrodynamic evolution models in which interaction effects break the conformal symmetry may lead to different initial conditions for the bulk viscous pressure, even in sign; this might be worth investigating in future studies.

In the present model, the initialization time for hydrodynamics is the free-streaming time τ_{fs} . It is common for model calculations to assume that this hydrodynamic initialization time is the same for all centralities and/or for different collision systems.⁹ There are, however, reasons to believe that systems with higher energy densities “hydrodynamize” faster [72]. To capture this effect we parametrize the free-streaming time τ_{fs} to include a dependence on the initially deposited transverse energy density $\bar{\epsilon}$ from Eq. (2):

$$\tau_{\text{fs}} = \tau_R \left(\frac{\langle \bar{\epsilon} \rangle}{\bar{\epsilon}_R} \right)^\alpha. \quad (19)$$

Here τ_R is a normalization factor for the duration of the free-streaming stage, and the parameter α controls its dependence on the magnitude of the “average” initial energy density in the transverse plane, defined by

$$\langle \bar{\epsilon} \rangle \equiv \frac{\int d^2x_\perp \bar{\epsilon}^2(\mathbf{x}_\perp)}{\int d^2x_\perp \bar{\epsilon}(\mathbf{x}_\perp)}. \quad (20)$$

We choose $\bar{\epsilon}_R = 4.0 \text{ GeV/fm}^2$ as an arbitrary reference scale; the resulting posteriors for α and τ_R depend on this choice.

B. Relativistic viscous hydrodynamics

After Landau matching, the energy-momentum tensor of the system is evolved with second-order relativistic viscous hydrodynamics [49–51, 73, 74]. In this work we use the dissipative fluid dynamics code MUSIC [75] and focus on the midrapidity

region where we can approximate the dynamics as effectively (2+1)-dimensional with longitudinal boost-invariance. Conservation of energy and momentum

$$\partial_\mu T^{\mu\nu} = 0, \quad (21)$$

$$T^{\mu\nu} = \epsilon u^\mu u^\nu - (p + \Pi) \Delta^{\mu\nu} + \pi^{\mu\nu} \quad (22)$$

provides evolution equations for the energy density and flow for given viscous flows. The latter (i.e. the shear stress $\pi^{\mu\nu}$ and bulk viscous pressure Π in Eq. (22)) follow their own relaxation equations:

$$\tau_\Pi \dot{\Pi} + \Pi = -\zeta\theta - \delta_{\Pi\Pi}\Pi\theta + \lambda_{\Pi\pi}\pi^{\mu\nu}\sigma_{\mu\nu}, \quad (23)$$

$$\begin{aligned} \tau_\pi \dot{\pi}^{\langle\mu\nu\rangle} + \pi^{\mu\nu} &= 2\eta\sigma^{\mu\nu} - \delta_{\pi\pi}\pi^{\mu\nu}\theta + \varphi_7\pi_\alpha^{\langle\mu} \pi^{\nu\rangle\alpha} \\ &\quad - \tau_{\pi\pi}\pi_\alpha^{\langle\mu} \sigma^{\nu\rangle\alpha} + \lambda_{\pi\Pi}\Pi\sigma^{\mu\nu}. \end{aligned} \quad (24)$$

Here $\dot{\Pi} = u^\lambda \partial_\lambda \Pi$, $\dot{\pi}^{\langle\mu\nu\rangle} = \Delta_{\alpha\beta}^{\mu\nu} u^\lambda \partial_\lambda \pi^{\alpha\beta}$, $\theta = \partial_\lambda u^\lambda$, and $\sigma^{\mu\nu} = \Delta_{\alpha\beta}^{\mu\nu} \partial^\alpha u^\beta$, with $\Delta_{\alpha\beta}^{\mu\nu}$ from Eq. (15).

The equilibrium properties of QCD matter are encoded in the equation of state which enters Eq. (22) through the pressure $p = p(\epsilon)$. The near-equilibrium dynamics of QCD matter is controlled by first and second-order transport coefficients that enter in Eqs. (23,24). The first-order transport coefficients are the shear and bulk viscosities, η and ζ . Second-order transport coefficients entering into our hydrodynamic equations are $\delta_{\Pi\Pi}$, $\lambda_{\Pi\pi}$, $\delta_{\pi\pi}$, φ_7 , $\tau_{\pi\pi}$, and $\lambda_{\pi\Pi}$, as well as the shear and bulk relaxation times τ_π and τ_Π .

For the equilibrium properties we use an equation of state matched to (i) a lattice calculation [2] at high temperatures and (ii) a hadron resonance gas at lower temperatures (see Refs. [53, 54] for details). The hadron content of the resonance gas is chosen to be consistent with that of the hadronic afterburner SMASH[57] used in this work. While this consistency is important, the matching procedure does carry some uncertainties (see. e.g., Ref. [76]) which are not explored in this work.

The shear and bulk viscosities, η and ζ , are parametrized as functions of temperature, and measurements are used to constrain them.¹⁰ They are discussed in more detail below.

With even less quantitative theoretical guidance available than for the first-order transport coefficients, the second-order ones should similarly be parametrized and constrained from measurements. Some studies¹¹ suggest, however, that the second-order transport coefficients have a smaller effect than first-order ones on the evolution of the plasma. In this work we therefore apply the same strategy as for the first-order transport coefficients η and ζ only to the shear relaxation time τ_π , whereas all other second-order transport coefficients are related to first-order ones by using parameter-free relations derived in a moment expansion of the Boltzmann equation [80].

¹⁰ In general, if conserved charges are taken into account (which is not done here), the transport coefficients also depend on chemical potentials.

¹¹ The effects of second-order transport coefficients on the hydrodynamic evolution were numerically checked in [77, 78] and found to be small. The hydrodynamic equations of motion used here are based on the Grad approach and constructed such that the second-order transport effects are assumed to be small compared to the first-order ones [79].

⁹ See, e.g., Ref. [71] for examples and exceptions.

All transport coefficients depend on the equilibrium properties of the system, which we characterize with the temperature T .¹⁰ We parametrize the ratios of shear and bulk viscosity to entropy density — the unitless specific viscosities — instead of parametrizing the viscosities themselves. A depiction of the parametrizations for the specific bulk and shear viscosities is shown in Fig. 1.




FIG. 1. Depictions of the parametrizations of specific bulk (left) and shear (right) viscosity as functions of temperature. The specific bulk viscosity has the form of a skewed Cauchy distribution, while the specific shear viscosity is piecewise-linear, with in general two different slopes. Both shear and bulk viscosities are required to be positive-definite to satisfy the second law of thermodynamics. The example for (η/s) shown here has a positive low-temperature and high-temperature slope ($a_{\text{low}}, a_{\text{high}} > 0$).

For the specific shear viscosity, η/s , we assume that it has a single inflection point at or above the deconfinement transition [81]. The position of this inflection point in temperature, T_η , is a parameter, as is the value of η/s at this point, $(\eta/s)_{\text{kink}}$. A linear dependence of η/s on temperature is assumed, with slopes a_{low} below and a_{high} above the inflection point, with both positive and negative slopes allowed. Negative values for η/s are not allowed. The formula for this parametrization is

$$\frac{\eta}{s}(T) = \max \left[\frac{\eta}{s} \Big|_{\text{lin}}(T), 0 \right], \quad (25)$$

with

$$\frac{\eta}{s} \Big|_{\text{lin}}(T) = a_{\text{low}}(T - T_\eta) \Theta(T_\eta - T) + (\eta/s)_{\text{kink}} + a_{\text{high}}(T - T_\eta) \Theta(T - T_\eta). \quad (26)$$

Theoretically expected is a negative slope at temperatures below T_η , i.e. $a_{\text{low}} < 0$, and a positive slope at temperature above T_η , i.e. $a_{\text{high}} > 0$ [82]. Nevertheless, in this work, we will allow both slopes to take negative and positive values: the aim is to ascertain whether phenomenological constraints are consistent with the theoretical expectations.

For the specific bulk viscosity, we assume that it peaks near the deconfinement temperature and that this single peak can be captured with a skewed Cauchy distribution:

$$\frac{\zeta}{s}(T) = \frac{(\zeta/s)_{\text{max}} \Lambda^2}{\Lambda^2 + (T - T_\zeta)^2}, \quad (27)$$

$$\Lambda = w_\zeta [1 + \lambda_\zeta \text{sign}(T - T_\zeta)].$$

Here T_ζ is the position and $(\zeta/s)_{\text{max}}$ the value of the peak; w_ζ and λ_ζ control the width and skewness of the Cauchy distribution, respectively. Allowing for a non-vanishing skewness is a generalization compared to Ref. [27].

Previous studies [7–10, 83] suggest that ζ/s for QCD peaks near the deconfinement transition. The functional form of its temperature-dependence is still not well understood. Below the transition ($T \lesssim 150$ MeV), the bulk viscosity is understood to be non-zero. We emphasize that we do not attempt to describe the dependence of the bulk viscosity below the particleization temperature of our model (discussed in the next section) which is never smaller than 135 MeV. The fact that our parametrization of $(\zeta/s)(T)$ rapidly approaches zero at low temperature should therefore not be read as a physical feature: this low temperature range is never described by the hydrodynamic code, but rather microscopically by a hadronic transport model. While we thus cannot make any statements about the bulk viscosity of QCD matter at these low temperatures it has recently been estimated in the SMASH transport model [7].

Previous theoretical work [80, 84–88] suggests that, in the absence of conserved charges, the shear relaxation time can be well captured by following temperature dependence:

$$T\tau_\pi(T) = b_\pi \frac{\eta}{s}(T) \quad (28)$$

where b_π is a constant that we consider unknown. The linearized causality bound [89] requires $b_\pi \geq (4/3)/(1 - c_s^2) \geq 2$. Refs. [80, 84–87] showed for a variety of weakly and strongly coupled theories other than QCD that this causality bound is respected, with b_π varying between ~ 2 and ~ 6 ; we use these values to limit the prior range explored for b_π in our parameter estimation.

Previous investigations of the effects of the shear relaxation time and other second-order transport coefficients on soft hadronic observables have found them to be of modest phenomenological importance [25, 77, 78, 90], consistent with general theoretical expectations (see e.g. Ref. [91]). Nevertheless varying the shear relaxation time in this work provides additional quantitative insights into the typical magnitude of effects from a second-order coefficient on the Bayesian constraints for the first-order transport coefficients.

C. Particleization

Particleization is not a physical process but a change of language from a description in terms of macroscopic fluid dynamical degrees of freedom to a microscopic kinetic description in terms of particles with positions and momenta. We here implement it on a surface of constant “switching” or “particleization” temperature T_{sw} . In principle, this translation requires simultaneous applicability of both approaches. Since hydrodynamics rapidly breaks down below the confinement transition because the mean-free path increases as a consequence of color neutralization, while the strongly-coupled nature of the color confinement process itself makes kinetic theory inapplicable during the phase transition, this condition puts rather tight theoretical constraints on the temperature range for the particleization procedure. We here impose these constraints through a

prior range within which we sample the particlization temperature. As we will see below, experimental data on hadronic yields provide strong constraints on T_{sw} , given the assumption that particlization happens at $\mu_i = b_i\mu_B + s_i\mu_S + q_i\mu_Q$ with $\mu_B = \mu_S = \mu_Q = 0$ (Eq. 31).

The Cooper-Frye [92, 93] prescription for particlization [94] is used to convert all the energy and momentum of the fluid into hadrons on the switching hypersurface Σ . The formula for the Lorentz-invariant particle momentum spectrum of particles of species i with degeneracy g_i in terms of their kinetic phase-space distribution $f_i(X; P)$ is given by

$$P^0 \frac{dN_i}{d^3P} = \frac{g_i}{(2\pi)^3} \int_{\Sigma} d^3\sigma_{\mu} P^{\mu} f_i(X; P). \quad (29)$$

The integral goes over the switching hypersurface Σ with normal vector $\sigma_{\mu}(X)$. The distribution function $f_i(X; P)$ must be chosen such that it reproduces the hydrodynamic energy-momentum tensor of the fluid on the particlization surface:

$$T^{\mu\nu}(X) = \sum_i g_i \int \frac{d^3P}{(2\pi)^3 P^0} P^{\mu} P^{\nu} f_i(X; P). \quad (30)$$

Without any hydrodynamic information on all the infinitely many other momentum moments of the distributions functions, and no hydrodynamic information on how to split the fluid $T^{\mu\nu}$ into contributions from different hadron species i as written in Eq. (30), this leaves infinitely many choices for the distribution functions $f_i(X, P)$. If the QGP fluid were an ideal fluid in perfect local kinetic and chemical equilibrium, the choice for $f_i(X, P)$ would be unambiguous: it would be of local equilibrium form [92, 93], with the local rest frame velocity provided by hydrodynamics and the temperature and chemical potentials fixed by the local-rest-frame energy density and the chemically equilibrated hadronic particle densities. In this case hadronic chemical potentials would be constrained by the equilibrium relations

$$\mu_i = b_i\mu_B + s_i\mu_S + q_i\mu_Q \quad (31)$$

where μ_B , μ_S and μ_Q are the chemical potentials associated with the conserved baryon number, strangeness and electric charge, and (b_i, s_i, q_i) are the baryon, strangeness and electric charges carried by hadron species i . All these chemical potentials are zero in this work, $\mu_B = \mu_S = \mu_Q = 0$, reflecting the approximately zero net baryon, strangeness and electric charge near midrapidity at top RHIC and LHC collision energies. As it is, the QGP is a dissipative fluid with nonzero dissipative flows contributing to $T^{\mu\nu}$ on the particlization surface, and since hydrodynamics does not provide any microscopic information on how the system evolved to this surface we are left with a large and irreducible ambiguity as to the choice of local momentum distributions for the different hadron species [95–98].

Working off the hypothesis that the color-confining hadronization process is very strongly coupled and involves a multitude of different possible color-neutralizing microscopic channels [99], and with strong support from experimental observations [100, 101] we will assume that the chemical potentials for the different hadronic species satisfy the above chemical equilibrium relations (31) at particlization as long as we

perform it soon after completion of hadronization. However, the dissipative flows in $T^{\mu\nu}$ reflect *deviations of the hadrons' momentum distributions and yields from local thermodynamic equilibrium*, arising from dissipative corrections. To specify these deviations one might want to argue that the distribution functions $f_i(X; P)$ should solve a set of coupled Boltzmann equations, but this begs the question what should be assumed for the form of the collision terms and the initial conditions (both of which must be expected to be strongly affected by the proximity in space and time of the preceding hadronization process about whose microscopic dynamics we know very little).

In such a situation of irreducible theoretical ambiguity it makes sense to ask what constraints the experimental data may provide. To pose this question we here consider four different models of viscous corrections to the local equilibrium distribution function:

1. Grad's method [102] (a.k.a. 14-moments method in the relativistic context [95, 103–108]);
2. the first-order Chapman-Enskog (CE) expansion in the Relaxation Time Approximation [109–111];
3. the Pratt-Torrieri-Bernhard (PTB) modified equilibrium distribution [53, 112];
4. the Pratt-Torrieri-McNelis (PTM) modified equilibrium distribution [55, 112].

Given the values of the 10 components of the energy-momentum tensor, these models are used to determine how energy and momentum are distributed among hadronic species and across momentum. By performing Bayesian parameter estimation, combined with Bayesian Model Selection techniques using these four models, we aim to estimate the theoretical uncertainty in the extraction of the transport coefficients resulting from the viscous corrections at particlization.

We briefly describe the four models individually. For a more in-depth review and comparison of these models we refer the reader to Ref. [55].

1. Linear viscous corrections: Grad & Chapman-Enskog

The Grad and Chapman-Enskog methods have both been used extensively in the literature to study particle production in heavy ion collisions. They give corrections which are linear in the dissipative currents $\pi^{\mu\nu}$ and Π . Both ansätze should only be valid when these corrections to the thermal equilibrium distribution are small. In practice this approximation is often pushed to the limit or even beyond. In the following we describe the Grad and Chapman-Enskog methods in turn. We then discuss regularization that is applied similarly to both approaches when large viscous corrections are encountered.

a. Grad (or 14-moments) approximation: What we refer to as ‘‘Grad’s method’’ assumes that the correction to the local equilibrium distribution function can be expanded in powers of hadronic momentum. Including only the terms relevant for a system without a conserved charge yields

$$\delta f_i = f_{\text{eq},i} \bar{f}_{\text{eq},i} c_{\mu\nu} P^{\mu} P^{\nu} \quad (32)$$

where $\bar{f}_{\text{eq},i} \equiv 1 - \Theta f_{\text{eq},i}$, and Θ is 1 for fermions and -1 for bosons. Assuming that the coefficients $c_{\mu\nu}$ are species-independent and requiring the Landau-matching conditions yields the following expression for the viscous correction in terms of the dissipative currents:

$$\delta f_i^{\text{Grad}} = f_{\text{eq},i} \bar{f}_{\text{eq},i} \left[\Pi (A_T m_i^2 + A_E (u \cdot P)^2) + A_\pi \pi^{\mu\nu} P_{\langle\mu} P_{\nu\rangle} \right]. \quad (33)$$

Here A_T , A_E , and A_π are combinations of thermodynamic moments of the equilibrium distribution described in Ref. [55], m_i is the mass of the hadron species i , $P_{\langle\mu} P_{\nu\rangle} = \Delta_{\mu\nu}^{\alpha\beta} P_\alpha P_\beta$, and $\Delta_{\mu\nu}^{\alpha\beta}$ is defined in Eq. (15).

b. Linearized Chapman-Enskog expansion in the relaxation time approximation (CE RTA): The Chapman-Enskog (CE) expansion is a method to solve the Boltzmann equation by expanding in Knudsen number, which is the dimensionless ratio of microscopic to macroscopic length scales in the system. Although this series can be written down for a more general collisional kernel we here use the relaxation-time approximation (RTA) [110, 113],

$$P^\mu \partial_\mu f = -\frac{u \cdot P}{\tau_{\text{rel}}} (f - f_{\text{eq}}), \quad (34)$$

where f_{eq} is the local equilibrium distribution function, and the relaxation time τ_{rel} is assumed to be species- and momentum-independent. Expanding the distribution function in a series around the local equilibrium distribution, called the Chapman-Enskog series, and keeping only the first-order correction one finds

$$f = f_{\text{eq}} - \frac{\tau_{\text{rel}}}{u \cdot P} P^\mu \partial_\mu f_{\text{eq}} + \mathcal{O}(\partial^2). \quad (35)$$

Using the zeroth order conservation laws to rewrite derivatives of the temperature and flow velocity, as well as the Navier-Stokes relations $\Pi = -\zeta\theta$ and $\pi^{\mu\nu} = 2\eta\sigma^{\mu\nu}$ we finally obtain

$$\delta f_i^{\text{CE RTA}} = f_{\text{eq},i} \bar{f}_{\text{eq},i} \left[\frac{\Pi}{\beta_\Pi} \left(\frac{(u \cdot P)\mathcal{F}}{T^2} - \frac{P \cdot \Delta \cdot P}{3(u \cdot P)T} \right) + \frac{\pi_{\mu\nu} P^{\langle\mu} P^{\nu\rangle}}{2\beta_\pi (u \cdot P)T} \right]. \quad (36)$$

Again we refer to [55] for the definitions of \mathcal{F} , β_π and β_Π .

c. Handling large viscous corrections: The Grad and Chapman-Enskog momentum distributions discussed above assume $|\delta f| \ll f_{\text{eq}}$. The viscous correction δf scales linearly with the shear stress $\pi^{\mu\nu}$ and the bulk viscous pressure Π . It also scales either quadratically or linearly with the hadron momentum P . There are thus values of $\pi^{\mu\nu}$ and Π for which $|\delta f| > f_{\text{eq}}$ even for moderate (thermal) momenta. Moreover, even for small values of $\pi^{\mu\nu}$ and Π , $|\delta f| > f_{\text{eq}}$ at sufficiently large momenta.

In hydrodynamic simulations of heavy-ion collisions it is thus not uncommon to encounter $|\delta f| > f_{\text{eq}}$ in certain phase-space regions. Even though these regions are usually small enough to not contribute significantly to experimental observables, from a practical point of view one does need to specify a hadronic momentum distribution even when $|\delta f| \sim f_{\text{eq}}$. This

is commonly achieved by regulating the Grad or Chapman-Enskog viscous corrections to prevent $|\delta f| > f_{\text{eq}}$. In this work this is achieved locally by setting

$$\delta f \rightarrow \text{sign}(\delta f) \min(f_{\text{eq}}, |\delta f|) \quad (37)$$

in every cell.

The need for regulation of the linearized viscous corrections has motivated models that attempt to resum the viscous corrections to all orders. We now discuss two such prescriptions.

2. Exponentiated viscous corrections: Pratt-Torrieri-McNelis and Pratt-Torrieri-Bernhard

The approaches described in this subsection rely on the development of positive definite ‘‘modified equilibrium’’ distributions [53, 55, 112] that include the effects of the viscous pressures in the argument of the exponential function that characterizes the equilibrium distribution.

a. Pratt-Torrieri-McNelis (PTM): The Pratt-Torrieri-McNelis (PTM) distribution [55, 112] is defined as follows:

$$f_{\text{PTM}} = \mathcal{Z} \left[\exp \left(\frac{\sqrt{|\mathbf{P}'|^2 + m^2}}{T + \beta_\Pi^{-1} \Pi \mathcal{F}} \right) + \Theta \right]^{-1}. \quad (38)$$

Here the spatial momentum components have been transformed as $P_i = A_{ij} P'_j$ where

$$A_{ij} \equiv \left(1 + \frac{\Pi}{3\beta_\Pi} \right) \delta_{ij} + \frac{\pi_{ij}}{2\beta_\pi}. \quad (39)$$

The PTM ansatz has the feature that expanding to first order in the dissipative currents yields the usual linear Chapman-Enskog viscous correction discussed above. The yield of each hadron is corrected from its equilibrium yield by a scaling factor \mathcal{Z} which depends on the bulk viscous pressure as well as the hadron mass as described in Ref. [55].

b. Pratt-Torrieri-Bernhard (PTB): The Pratt-Torrieri-Bernhard (PTB) distribution [53, 112] is defined by

$$f_{\text{PTB}} = \frac{\mathcal{Z}_\Pi}{\det \Lambda} \left[\exp \left(\frac{\sqrt{|\mathbf{P}'|^2 + m^2}}{T} \right) + \Theta \right]^{-1}, \quad (40)$$

where \mathcal{Z}_Π is a scaling factor described in Ref. [53, 55] which again depends on the bulk viscous pressure but is species-independent. Λ is a momentum-transformation matrix operating on the spatial momentum components as $P_i = \Lambda_{ij} P'_j$ with

$$\Lambda_{ij} \equiv (1 + \lambda_\Pi) \delta_{ij} + \frac{\pi_{ij}}{2\beta_\pi}. \quad (41)$$

In particular, $\lambda_\Pi \neq \Pi/(3\beta_\Pi)$; instead, it is adjusted such that the total pressure of the system is matched. This method parametrizes the effect of the bulk viscous pressure on the particle yields and momentum spectra, and it does not reduce to the linear Chapman-Enskog correction in the limit of small Π .

The PTB distribution was used in the recent Bayesian analysis of Ref. [27]. It should be noted that, in contrast to the (unregulated) linearized Grad and Chapman-Enskog distributions, for both PTB and PTM distributions the matching constraint (30) is not satisfied exactly when the viscous stresses are large [55]. The slight matching inconsistencies introduced by the different regulation schemes discussed above were quantitatively studied in [55] and found to be acceptable in practice. For other approaches to regulate the viscous corrections to the distribution functions during particlization we refer the interested reader to Refs. [114–122].

We remind the reader that in the presence of bulk viscous stresses the bulk viscous corrections to the distribution function f_i shift the chemical equilibrium yields of the hadron species i [107] and thereby have the potential to affect the particlization temperature T_{sw} at which the hadron yields are consistent with the equilibrium relations (31) (see discussion in the following subsection).

D. Hadronic transport

In our hybrid model we transition to microscopic hadronic Boltzmann dynamics, simulated with the kinetic evolution code SMASH [57, 58], by imposing particlization at the switching temperature T_{sw} as described above. After particlization of the fluid, the resulting hadrons are allowed to scatter, form resonances, and decay. SMASH solves a tower of coupled Boltzmann equations for a system of hadronic resonances:

$$P^\mu \partial_\mu f_i(x; P) = C[f_i], \quad (42)$$

where f_i is the distribution function for hadronic species i and $C[f_i]$ is the collision term describing all scattering, resonance formation, and decays involving particle species i .

Past phenomenological studies [35–42] have found that including a hadronic afterburner improves the ability of a hydrodynamic model to describe the spectra of heavier hadronic states, such as protons. This transport approach allows different species to reach chemical and kinetic freezeout dynamically. This contrasts with other approaches where chemical and kinetic freezeout are enforced at specific temperatures.¹² At particlization, the momentum distributions and particle yields already deviate from their equilibrium relations at that temperature due to shear and bulk viscous effects. After switching to the afterburner, they continue to evolve until yields and momentum distributions cease changing. Most hadronic yields vary by less than 20% as a consequence of inelastic collisions in the afterburner phase, and the particlization temperature T_{sw} is therefore sometimes associated with a

chemical freeze-out temperature [38]. However, baryon and anti-baryon yields may change more significantly, due to the large annihilation cross section [34, 124]. As a result, this is not a precise relationship and we indeed find somewhat lower values for T_{sw} than the canonical chemical freeze-out temperatures extracted from static thermal model fits such as those in Refs. [100, 101].

We note that none of the parameters in the SMASH afterburner are varied in this work. We did validate, however, that the afterburner used in this work (SMASH) agrees well with the popular UrQMD implementation that has been used for decades. This comparison is discussed in Appendix H 2.

Treatment of the σ meson: At particlization the hydrodynamic energy-momentum tensor is converted into hadrons assuming the system has the thermodynamic properties of a hadron resonance gas. Though the σ meson can be formed as a resonance in the $\pi + \pi$ scattering channel, it has been shown in Ref. [125] that the contribution to the partition function from σ meson exchange (an isoscalar-scalar channel) is almost perfectly canceled by the repulsive isotensor-scalar channel in $\pi + \pi$ scattering. Based on this observation, it is generally agreed that the σ meson should be omitted from isospin-averaged hadron resonance gas models [125]. This is the approach we use in this work: the σ meson is *not* sampled at particlization, and correspondingly it is also omitted in the construction of the equation of state in the hadronic phase.¹³ In the hadronic afterburner, we still allow SMASH to dynamically form and decay σ resonances because they are an essential ingredient in explaining the $\pi + \pi$ cross section in SMASH. We note for reference that the Bayesian analysis in Ref. [27] did include the σ meson in both the sampling at particlization and the construction of the hadronic equation of state, making this one of its differences with the current analysis.

This concludes the discussion of the dynamical evolution model used in this work. We now proceed to a discussion of the statistical tools used in its calibration with experimental data.

IV. SPECIFYING PRIOR KNOWLEDGE

Before using a set of measurements to perform Bayesian inference, one must quantify the current state of knowledge. If we want to infer the likely values of model parameters, given some observed data, then we must quantify our belief about the model parameters before we see the data. If we want to compare models, given some observed data, we must quantify the likely values of each model’s parameters, as well as our belief in each model, before seeing the data. Broadly, before we use Bayesian inference to address a question in light of some observed data, our “prior” encodes our current state of knowledge before we have seen the selfsame data.

¹² For example, the partial chemical equilibrium approach [123] enforces chemical freezeout at a given temperature in ideal hydrodynamics, by introducing chemical potentials to conserve all hadronic multiplicities to a chosen chemical freeze-out values. This was a popular procedure before the widespread availability of hybrid codes (see e.g. [38, 39] for comparisons of these two approaches).

¹³ More details about the construction of the equation of state are provided in Appendix H 6. The physical effects on observables from excluding the σ meson from the hadron gas are studied in Appendix H 4.

Norm. Pb-Pb 2.76 TeV	$N[2.76 \text{ TeV}]$	[10, 20]	temperature of (η/s) kink	T_η	[0.13, 0.3] GeV
Norm. Au-Au 200 GeV	$N[0.2 \text{ TeV}]$	[3, 10]	(η/s) at kink	$(\eta/s)_{\text{kink}}$	[0.01, 0.2]
generalized mean	p	[-0.7, 0.7]	low temp. slope of (η/s)	a_{low}	[-2, 1] GeV ⁻¹
nucleon width	w	[0.5, 1.5] fm	high temp. slope of (η/s)	a_{high}	[-1, 2] GeV ⁻¹
min. dist. btw. nucleons	d_{min}^3	[0, 1.7 ³] fm ³	shear relaxation time factor	b_π	[2, 8]
multiplicity fluctuation	σ_k	[0.3, 2.0]	maximum of (ζ/s)	$(\zeta/s)_{\text{max}}$	[0.01, 0.25]
free-streaming time scale	τ_R	[0.3, 2.0] fm/c	temperature of (ζ/s) peak	T_ζ	[0.12, 0.3] GeV
free-streaming energy dep.	α	[-0.3, 0.3]	width of (ζ/s) peak	w_ζ	[0.025, 0.15] GeV
particlization temperature	T_{sw}	[0.135, 0.165] GeV	asymmetry of (ζ/s) peak	λ_ζ	[-0.8, 0.8]

TABLE I. A list of all priors used (see Sec. III for the definitions of the model parameters). All prior distributions are assumed to be uniform and nonzero within the range quoted, and zero outside. The Table does not exhibit the step functions that ensure non-negativity of the shear viscosity at all temperatures (see Eq. (25)).

It is essential that the construction of the prior distribution should not be informed by the same data that will be used in performing parameter estimation. In particular, the posterior of earlier analyses that used the same data sets should not in any way be used as a prior for a new analysis: it would be an attempt to use the same measurements twice, as well as being likely inconsistent given differences in the models.

When selecting a prior different factors must be considered. Theoretical constraints are important, including self-consistency concerns for the model and/or conservation laws or symmetries that must be respected, all of which are problem-specific. Within these constraints, a range of different priors is possible. There are methods aimed at reducing subjectivity in the choice of priors; examples include maximum-entropy priors [126]. In this work we focus on a careful selection of the range of the priors, rather than on the exact form of the prior probability distribution for each parameter. In the following two paragraphs we illustrate this selection process for a subset of the parameters.

a. Initial conditions: There are physical constraints on the prior for the width parameter w in T_RENTo: when choosing a reasonable range of values one must keep in mind that the electric charge radius of the proton is about 0.9 fm. The width parameter w in T_RENTo should likely not be allowed to be much smaller or larger than this value.

b. Transport coefficients: The ranges of allowed shear and bulk viscosities are important priors. First, both viscosities should be non-negative, to ensure the second law of thermodynamics, i.e. a positive entropy production rate. At the opposite end of the allowed range, the applicability of hydrodynamics becomes debatable when the viscous part of the energy-momentum tensor is large compared to the ideal part. This can happen when the shear and bulk viscosities are large. For self-consistency it is thus desirable that the prior ranges of η/s and ζ/s exclude unreasonably large values. Though exploring large values of the viscosities may be physically interesting, it would push the hydrodynamic component of our model outside its regime of validity. If the experimental data require larger viscosities than included in our prior, this should be visible in our posterior distributions for these transport coefficients, as well as in a low value for the model evidence, i.e. a bad fit of the data for any viscosity within the allowed prior range. The shear relaxation time also has physical constraints. As discussed in Section III B, there is a minimum value for $b_\pi = T\tau_\pi/(\eta/s)$ set by causality of the linearized equations,

$b_\pi \gtrsim 2$, and theoretical evaluations of b_π within a number of microscopic theories ranging from weakly to strongly coupled provide some theoretical guidance for the most likely range of this parameter.

In the present analysis, for simplicity all of the parameters (denoted by the vector \mathbf{x}) are assigned a uniform prior probability density $\mathcal{P}(\mathbf{x})$ on a finite range. These ranges are listed in Table I; as discussed above, they have been chosen with certain theoretical biases. The priors for different parameters are assumed to be independent, so that the joint prior is simply given by their product,

$$\mathcal{P}(\mathbf{x}) \propto \prod_i \Theta(x_i - x_{i,\text{min}})\Theta(x_{i,\text{max}} - x_i), \quad (43)$$

where i runs over all the model parameters in \mathbf{x} . Note that uniform priors are not uninformative priors. Moreover, the choice of priors in principle affects the results of the Bayesian parameter estimation, especially in situations where the data do not have sufficient information to correct prior prejudice. For instance, in this work, we require $(\eta/s)(T)$ and $(\zeta/s)(T)$ to be given by specific parametrizations, with each of the parameters sampled from a uniform prior. The resulting prior for $(\eta/s)(T)$ is, however, not uniform as a function of temperature; thus, our choice of parametrization informs our prior. A plot showing credible intervals for the prior for the shear and bulk viscosities is shown in Fig. 2. We see that this prior




FIG. 2. Credible intervals of the prior probability density for the specific bulk (left) and shear (right) viscosities that we use when performing Bayesian parameter estimation. The 60%, 90% and 100% credible intervals (C.I.) are shown.

encapsulates our belief that the bulk viscosity should have a

peak somewhere near the deconfinement transition temperature, and that the specific shear viscosity reaches a minimum in that region.

Nevertheless, we used a broad prior for η/s , allowing it to take either a maximum or a minimum in the deconfinement region. By doing so we tried to limit the theoretical bias of our prior for $(\eta/s)(T)$. When selecting the priors for the remaining model parameters we followed similar considerations, with the goal of ensuring that our posterior parameter constraints will be guided as much as possible by the heavy-ion data and not by prior prejudice.

It is important, however, to understand that in practice theoretical bias can never be fully avoided. In many cases it can be helpful or even needed: if highly constraining data are lacking, exploring the reaction of the posterior distribution to different prior theoretical assumptions can yield useful insights into the variability and reliability of model predictions and suggest strategies for decreasing this variability with targeted theoretical efforts. The Bayesian methodology accepts the reality of theoretical bias, but at the same time accounts for it quantitatively in the posterior probability distribution for the model parameters x . Sensitivity to our prior assumptions is further explored in Sec. IX.

V. BAYESIAN PARAMETER ESTIMATION WITH A STATISTICAL EMULATOR

In this section, we describe the statistical methodology used to estimate the likely model parameters by comparison with the experimental data. This non-trivial problem is tackled by the application of Bayesian parameter estimation with a physical model surrogate or “emulator”. The use of a model emulator is necessary when the physical simulation is computationally intensive. Performing a Bayesian parameter estimation requires evaluating the model’s prediction on arbitrary points in the relevant region of the parameter space. In theory, this involves evaluating the model for millions of different values for initial conditions, viscosities and other model parameters. Given a model that requires nearly $\mathcal{O}(10^3)$ CPU-hours to run a single set of input parameters, this would quickly become intractable as a method of exploring the posterior. The model emulator is designed precisely to solve this problem. The emulator can be understood as a computationally fast interpolation of the physical model simulation with an estimate of the interpolation uncertainty. The model is evaluated on a sample set of points in the parameter space, and the model’s predictions at these points are used to infer the predictions at other points in parameter space through the use of an emulator. Such an emulator dramatically reduces the numerical cost of estimating the posterior; the trade-offs are that it assumes a certain smoothness in the behavior of the model outputs, as well as introduces an additional source of uncertainty in the analysis: emulator uncertainty. In addition to describing Bayesian parameter estimation in general, we also discuss specifically the role of the emulator. Our discussion in this section builds upon Refs. [22–27, 67] where many of these techniques were previously applied to Bayesian parameter estimation in relativistic

heavy-ion physics. Additional information on Bayesian inference can be found in e.g. Ref. [33].

A. Overview of Bayesian Parameter Estimation

Bayesian parameter estimation is a systematic approach to infer the probability distribution of model parameters (x) by comparing theoretical calculations (y_x) to experimental data (y_{exp}). The starting point is the prior distribution $\mathcal{P}(x)$ that encodes the current state of knowledge regarding the model parameters x before making comparison with data (Section IV). The posterior distribution of model parameters after model-to-data comparison, $\mathcal{P}(x|y_{\text{exp}})$, is given by Bayes’ theorem,

$$\mathcal{P}(x|y_{\text{exp}}) = \frac{\mathcal{P}(y_{\text{exp}}|x)\mathcal{P}(x)}{\mathcal{P}(y_{\text{exp}})}, \quad (44)$$

where $\mathcal{P}(y_{\text{exp}}|x)$ is the “likelihood” that the model agrees with experimental measurement, given the parameters x , and the normalization $\mathcal{P}(y_{\text{exp}})$ is called the “Bayesian evidence”. The exact form of the likelihood is often unknown, as it depends on the probability distribution of the experimental and theoretical uncertainties. In this work, we follow the common assumption that the likelihood can be taken to be a multivariate normal distribution. This choice is justified when uncertainties are normally distributed. With this choice of likelihood function, the logarithm of $\mathcal{P}(y_{\text{exp}}|x)$ contains the quadratic form of the difference between the measurement and the prediction $\Delta y_x = y_x - y_{\text{exp}}$,

$$\log [\mathcal{P}(y_{\text{exp}}|x)] = -\frac{1}{2} \ln [(2\pi)^n \det \Sigma] - \frac{1}{2} \Delta y_x^T \Sigma^{-1} \Delta y_x.$$

Here, n is the number of observation points (i.e. the length of the vector y_{exp}), and Σ is a covariance matrix that encodes both experimental and model uncertainties, as well as correlations among uncertainties. These correlations are generally not readily available experimentally. As such the treatment of uncertainties can become a relatively complex question. We discuss the treatment of uncertainties and the covariance matrix separately in Section VC below.

In principle, in order to calculate the posterior, one is faced with the task of calculating the evidence $\mathcal{P}(y_{\text{exp}})$. For many problems of interest the required high dimensional integration can be numerically challenging or even intractable. Fortunately, when performing Bayesian parameter estimation, knowledge of the relative probability of different points in parameter space is sufficiently interesting in itself. That is, as the evidence $\mathcal{P}(y_{\text{exp}})$ does not depend on the parameters x , it is sufficient to consider the proportionality

$$\mathcal{P}(x|y_{\text{exp}}) \propto \mathcal{P}(y_{\text{exp}}|x)\mathcal{P}(x).$$

Methods for estimating the posterior which take advantage of this include Markov Chain Monte Carlo. Therefore, when we discuss or plot the posterior of parameter estimates throughout this section, we implicitly mean the unnormalized posterior. Hence, we are interested in the relative probability density of each parameter set, and *not* the absolute probability.

Because the plotted posterior for the model parameters in general does not contain information about this normalization, it is imperative to check the level of agreement between the posterior prediction of observables to assess quantitatively how well the model can describe the experimental data. As was explained in Section II, it is meaningless to ponder on the posterior estimates of parameters for a model which poorly explains the observed data. Thus, in Section VIII C, we will also explore how well the model observables sampled from the posterior fit the experimental data. An estimation of the evidence $\mathcal{P}(\mathbf{y}_{\text{exp}})$ becomes necessary if we want to compare models in a Bayesian framework and this will be discussed in Section XI.

Simultaneous constraints from multiple collision systems: When combining constraints from different experiments, RHIC and LHC for example, the joint likelihood function is assumed to be the product of the individual likelihoods for each system:

$$\mathcal{P}(\mathbf{y}_{\text{exp}}^{\text{LHC}}, \mathbf{y}_{\text{exp}}^{\text{RHIC}} | \mathbf{x}) = \mathcal{P}(\mathbf{y}_{\text{exp}}^{\text{LHC}} | \mathbf{x}) \mathcal{P}(\mathbf{y}_{\text{exp}}^{\text{RHIC}} | \mathbf{x}). \quad (45)$$

The parameter values that maximize the joint likelihood strike a compromise between maximizing the individual likelihoods.

Importantly, one must decide which parameters are shared for the different collision systems. Naively, one could expect that all parameters should be shared; in reality this depends partly on how the model parameters were defined.

We highlight that comparisons with measurements can always help determine if model assumptions need to be relaxed. If RHIC and LHC measurements could be described independently by the model but not simultaneously, it could be an indication that the $\sqrt{s_{\text{NN}}}$ dependence of certain parameters needs to be revisited, i.e., that it may not be correct to enforce the same value of certain parameters at RHIC and the LHC. We will compare more complex models which relax some of these assumptions by estimating Bayes factors in Section XI D.

Inclusion of data at two very different collision energies raises the question where and how we make allowance for $\sqrt{s_{\text{NN}}}$ dependence of the model parameters for which we interrogate the data for constraints. Answers are provided in the following paragraphs:

a. Initial stage model: Because T_{RENT}o is a parametric initial condition model, not a dynamical one, many of its parameters should, in principle, be beam-energy dependent.¹⁴ Generically, we assume that at high collision energies the parameters that we try to extract from experiment evolve sufficiently slowly with $\sqrt{s_{\text{NN}}}$ that their change from RHIC to LHC can be ignored. As an exception we retain the $\sqrt{s_{\text{NN}}}$ dependence of the normalization N of the energy density in T_{RENT}o, because it is directly responsible in our model for the large increase of mid-rapidity particle and energy production from RHIC to LHC. Rather than parametrizing its $\sqrt{s_{\text{NN}}}$

dependence, we simply use two independent normalizations at $\sqrt{s_{\text{NN}}} = 200$ and 2760 GeV, labeled by $N[0.2 \text{ TeV}]$ and $N[2.76 \text{ TeV}]$, respectively. We also point out that in Section XI D we use Bayesian Model Selection to explore whether experimental data would prefer a dependence of the nucleon width w in T_{RENT}o on $\sqrt{s_{\text{NN}}}$. The free-streaming time (19) is allowed to depend on $\sqrt{s_{\text{NN}}}$ implicitly, through the deposited energy density.

b. Transport coefficients: The specific shear and bulk viscosities, as well as the second-order transport coefficients in our hydrodynamic approach, are medium properties that (for systems without conserved charges) depend only on the temperature of the plasma. Their parametrizations as functions of temperature, $(\eta/s)(T)$ and $(\zeta/s)(T)$, are therefore independent of $\sqrt{s_{\text{NN}}}$.

c. Particlization: We use the same particlization temperature T_{sw} at RHIC and at the LHC. As discussed in Secs. III C and III D, particlization is assumed to happen at T_{sw} with chemical potentials (31) that, in a static environment, reflect chemical equilibrium relations between the hadronic yields at T_{sw} . It is known that, when initialized in this way at T_{sw} , hadronic transport changes these yield ratios much more slowly than they would in a chemically equilibrated fluid [34, 38, 42, 99, 128]. The finally observed hadronic chemical composition is therefore largely set at particlization [128], linking the switching temperature T_{sw} to the chemical decoupling temperature.¹⁵ At both collision energies matter with approximately zero baryon chemical potential is produced which passes through the hadronization phase transition at the same temperature. Based on strong experimental and theoretical evidence [99–101, 128] we expect the chemical equilibrium relations (31) between the chemical potentials of the different hadronic species to be broken soon after hadronization is complete, which should occur at the same value for T_{sw} at both collision energies. We consider this a combined theoretical and empirical prior for T_{sw} .

B. Physical model emulator

Throughout this study, we define an emulator as a map from a point in the multidimensional parameter space to the mean vector and covariance matrix of the distribution of all the predicted model observables of interest. This map provides a non-parametric estimation of the physical model calculations at arbitrary points in the region of the parameter space of interest. It is “non-parametric” because predictions at novel parameter points are not made by constructing explicit functional relations between model predictions and parameter inputs, but

¹⁴ For example, in the color glass condensate effective theory for QCD at very high energies, the only relevant scale is the saturation scale Q_s , which controls correlations in the transverse direction and which runs with the energy of the collision system [127]. This suggests that the nucleon width in T_{RENT}o should perhaps have a similar $\sqrt{s_{\text{NN}}}$ dependence.

¹⁵ Let us emphasize again that, for the same values of T_{sw} and μ_i (31), bulk viscous pressure corrections (which are different at RHIC and LHC, due to different expansion rates) will lead to different hadron yields dN_i/dy in collisions at RHIC and LHC energies. To the best of our knowledge this dissipative correction to thermal equilibrium model fits of hadron yield ratios has not been previously studied, but it is automatically taken into account in our Bayesian inference analysis.

rather are obtained by modeling the way that predictions at any point are correlated with known calculations at other parameter points. This sample of points in parameter space where we know the physical model calculations are called the design points ($\mathbf{x}_i; i = 1, \dots, m$).

The parameter design samples are chosen carefully using the Latin hypercube sampling technique, which randomly fills the volume of parameter space yielding a uniform distribution for each parameter while maximizing the distance between adjacent points. For models with sufficient smoothness, the number of design points necessary to achieve a certain level of prediction accuracy scales linearly with the dimension of the parameter space¹⁶ [129]. In this work we have taken a Latin hypercube design of 500 points. At each design point, the full model is run 2500 times for each collision system (Au+Au at RHIC or Pb+Pb at LHC), each time with a randomly fluctuating initial condition. The 2500 events are then ordered according to the yield of charged hadrons in each event $dN_{\text{ch}}/d\eta$ to define centrality classes, and observables are averaged over the fluctuating events in centrality bins which match those given by the experimental measurements.

With the training data available, the following steps define the construction and training of the emulator.

a. Dimensionality reduction via Principal Component Analysis: When comparing the model output with experimental data, we are faced with the large dimensionality of the output. Many of the model observables carry correlated information. As a simple example, an increased normalization of initial energy density increases pion multiplicity in all centrality bins. Therefore, the predicted value of multiplicity at different centralities effectively equals a single degree-of-freedom in response to the change of the normalization parameter. To put it another way, a small subspace of the full model output carries nearly all of the information about the model parameters. Therefore, we apply principal component analysis as a dimensionality reduction method. Suppose an array of observations y_i ($i = 1, \dots, n$) are calculated at each of the $m = 500$ design points j . They are organized as an $n \times m$ matrix Y with elements y_{ij} . First, for each of the observables y_i , we compute its mean μ_i and standard deviation σ_i over the sample of m design points. Then, each of the n observables is standardized by subtracting the mean and dividing by the standard deviation, yielding an $n \times m$ matrix \tilde{Y} with elements $\tilde{y}_{ij} = (y_{ij} - \mu_j)/\sigma_j$ for $j = 1, \dots, m$. Secondly, we define a new set of ‘‘observables’’ z_i which are linear combinations of the standardized observables: $z_i = O_{ik}\tilde{y}_k$. One seeks an optimized set of z_i such that the linear correlations between different z -observables vanish:

$$\begin{aligned} \langle \delta z_i \delta z_j \rangle &= \frac{1}{m} \sum_{k=1}^m (O\tilde{Y})_{ik} (O\tilde{Y})_{jk} = \frac{1}{m} (O\tilde{Y}\tilde{Y}^T O^T)_{ij} \\ &= \lambda_i \delta_{ij} \equiv \text{diag}\{\lambda_1, \dots, \lambda_n\}, \end{aligned} \quad (46)$$

where δz_i denotes the deviation of the z_i from their mean. Therefore, the coefficients O_{ij} that define z_i are simply the el-

ements of the orthogonal matrix that diagonalizes the covariance matrix of \tilde{y}_i . This optimized set of z_i are the so-called principal components. The rows of O are organized such that the eigenvalues λ_i , which are the variances of the z_i , have a descending order. In this way, each successive principal component explains less variance in the standardized observables. This allows us to reduce the standardized observable space to a much smaller subspace, which captures most of the information about the parameters. One should remember that the principal component analysis can only remove linear correlations among observables, so it is important to check that there are no significant non-linear correlations. This is demonstrated in Appendix C.

In our experience, a very small fraction of the total number of principal components is generally sufficient to capture most of the model observables’ dependence on the parameters. This follows from the strong linear correlations present in many pairs of observables. Pairs of observables with stronger linear correlations carry less mutual information about the parameters; knowledge of one observable is nearly sufficient to know the value of the other. Gaussian processes are only trained on this subset of dominant principal components.

From a practical point of view, we also observed that it is important not to include too many principal components: this limits the risk that the emulator overfits the noise present in the simulation data.

b. Interpolating each Principal Component by a Gaussian Process: Each dominant principal component is interpolated with a unique Gaussian process. The spirit of a Gaussian process regressor is to infer the outputs of the target (scalar) function $y = M(x)$ ¹⁷ by a distribution of functions denoted by \mathcal{GP} : $f(x) \sim \mathcal{GP}(\text{mean}(x), \text{cov}(x, x'))$. This distribution is a multivariate normal distribution specified by a mean $\mu(x)$ and a covariance $\text{cov}(x, x')$, so that the expectation value of the output at a given x is

$$\langle f(x) \rangle = \text{mean}(x), \quad (47)$$

and the correlation of the output between two independent inputs x, x' is

$$\langle \delta f(x) \delta f(x') \rangle = \text{cov}(x, x'), \quad (48)$$

where $\delta f(x) = f(x) - \text{mean}(x)$.

To find the desired distribution of functions that emulates $M(x)$, one starts with a distribution that is completely agnostic to the target function $M(x)$. In this study this distribution, referred to as the unconditioned Gaussian process, is assumed to have mean $\mu(x) = 0$ ¹⁸ and a covariance function $k(x, x')$ (the so-called kernel function). A Gaussian process makes a prediction at m_* novel inputs X_* according to the correlations with known values that have been calculated at the m training

¹⁶ This scaling of interpolation uncertainty with design size is explored in [31] for a different set of observables.

¹⁷ In this context, the output of the target function is one of the dominant principal components.

¹⁸ It can happen that near the boundaries of parameter space the model prediction for some principal component is nonzero. In this case it may be beneficial to include a non-zero mean function in the Gaussian Process. We do not explore this in this work.

inputs X . Consistency requires that the joint distribution of outputs at both training and novel inputs is also multivariate normal with zero mean,

$$\begin{bmatrix} \mathbf{f}(X) \\ \mathbf{f}(X_*) \end{bmatrix} \sim \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} K(X, X) & K(X, X_*) \\ K(X_*, X) & K(X_*, X_*) \end{bmatrix}\right) \quad (49)$$

where $K(X, X_*)$ is the $m \times m_*$ matrix whose elements are composed of the pointwise covariances $k(\mathbf{x}_p, \mathbf{x}_q)$ between pairs of training points \mathbf{x}_p and prediction points \mathbf{x}_q . Then, one conditions the random vector $\mathbf{f}(X)$ on the training outputs $M(X)$ to obtain the probability distribution of $\mathbf{f}(X_*)$ given training data. The mean and covariance can be obtained by the properties of the multivariate normal distribution,

$$\mathbf{f}(X_*) \sim \mathcal{GP}(\text{mean}(X_*), \text{cov}(X_*, X_*)) \quad (50)$$

$$\text{mean}(X_*) = K(X_*, X) [K(X, X)]^{-1} M(X), \quad (51)$$

$$\begin{aligned} \text{cov}(X_*, X_*) &= K(X_*, X_*) \\ &\quad - K(X_*, X) [K(X, X)]^{-1} K(X, X_*). \end{aligned} \quad (52)$$

Focusing on a single novel input, the prediction with uncertainty quantification of the target function is $M(x_*) \approx \text{mean}(x_*) \pm \sqrt{\text{cov}(x_*, x_*)}$. Eqs. (51), (52) imply that if x_* coincides with one of the training inputs then the mean agrees with the training output with vanishing uncertainty.

In addition to the training data, choosing the kernel function $k(x, x')$ is another key step in Gaussian process regression. An independent kernel function $k(\mathbf{x}_p, \mathbf{x}_q)$ is assigned to each dominant principal component, and is given by the sum of a squared-exponential kernel $k_{\text{exp}}(\mathbf{x}_p, \mathbf{x}_q)$ and white-noise kernel $k_{\text{noise}}(\mathbf{x}_p, \mathbf{x}_q)$,

$$k(\mathbf{x}_p, \mathbf{x}_q) = k_{\text{exp}}(\mathbf{x}_p, \mathbf{x}_q) + k_{\text{noise}}(\mathbf{x}_p, \mathbf{x}_q). \quad (53)$$

The squared-exponential kernel is given by

$$k_{\text{exp}}(\mathbf{x}_p, \mathbf{x}_q) = C^2 \exp\left(-\frac{1}{2} \sum_{i=1}^s \frac{|x_{p,i} - x_{q,i}|^2}{l_i^2}\right) \quad (54)$$

where C^2 is the unknown auto-correlation hyperparameter. The index i runs over all s parameters, and each parameter is assigned an uncertain hyperparameter l_i . This length-scale l_i controls the smoothness of the response of the principal component output to a change in the i^{th} parameter. The white-noise kernel is given by

$$k_{\text{noise}}(\mathbf{x}_p, \mathbf{x}_q) = \sigma_{\text{noise}}^2 \delta_{p,q} \quad (55)$$

where $\delta_{p,q}$ is the Kronecker delta, while σ_{noise} is an uncertain hyperparameter controlling the amount of statistical spread present in the principal component. The σ_{noise} is present because our model calculations average over a finite number of initial conditions and a finite number of particles.

All of the hyperparameters C, l_i and σ_{noise} are assigned a possible window, and then simultaneously optimized inside this window such that they maximize the likelihood of fit of the Gaussian process to the training calculations. This likelihood includes a complexity penalty, to reduce the potential for overfitting. This procedure is automated, and performing emulator validation is necessary to check that each kernel function has hyperparameters which are not underfit or overfit [130].

c. Reconstructing the observables: The predictions for principal components are then grouped and transformed back into the observables via the inverse PCA transformation. Variances of those non-dominant principal components on which we did not train Gaussian processes are included as prediction uncertainty. These neglected principal components in fact behave similarly to noise terms. We use this feature to actually replace them by white noise (variance which is uncorrelated point-to-point in parameter space) terms as an estimation of their contributed uncertainty.

A more detailed description of the above procedure can be found in [53]. We note that our use of transverse-momentum-integrated observables, principal component analysis, and Gaussian process model emulator for performing Bayesian parameter estimation for heavy-ion collisions is very similar to those put forward in the seminal study [23].

C. Treatment of uncertainties

We divide our uncertainties into three different sources: experimental uncertainties, interpolation and statistical model uncertainties, as well as systematic model discrepancies.

a. Experimental uncertainties: In general, experimental collaborations do not report the error covariance matrix between different observables, or for different momentum bins of the same observable. As such, we only have access to the systematic uncertainties of individual observables or bins, with limited or no information on possible correlations. Assuming no correlations among the errors associated with the n observables results in a diagonal covariance matrix for the experimental systematic covariance:

$$\Sigma_{\text{sys}} = \text{diag}(\sigma_{\text{sys},1}^2, \dots, \sigma_{\text{sys},n}^2). \quad (56)$$

In principle, the systematic uncertainties have nonzero correlations. Although it is possible to use our model to estimate correlations between the mean values of different observables, we do not have information about other sources of systematic uncertainties. Without knowledge of the experimental covariance matrix we can only make assumptions regarding the form and magnitude of the correlations. We have tested the effect of this approach on the parameter posteriors in Appendix D. However, we did not use this approach in general in the body of this work.

b. Interpolation and statistical model uncertainties: The statistical uncertainty which is present in our model calculations results primarily from averaging over a finite number of fluctuating initial conditions, and to a lesser extent sampling a finite number of particles during particlization. These result in a statistical spread in each of the principal components (recall from Section VB that it is the principal components that are interpolated, not the individual observables).

The total interpolation uncertainty is

$$\Sigma_{\text{interp}} = \Sigma_{\text{trunc.}} + \Sigma_{\text{GP}}. \quad (57)$$

The covariance Σ_{GP} contains the total covariance of all the Gaussian Processes (one for each dominant principal compo-

ment), including both interpolation and statistical uncertainties. The covariance $\Sigma_{\text{trunc.}}$ contains the total covariance of all the remaining principal components to which Gaussian processes were not fit and which were replaced by noise terms.

c. Additional systematic model discrepancy: Our model of heavy ion collisions is unavoidably imperfect. Therefore there exist additional sources of systematic discrepancy in our model when we use it to describe physical observations. Quantifying and interpreting the associated uncertainties presents a challenging problem [131].

Some of these uncertainties could in principle be quantified, although for practical reasons we will not attempt to do so in this work. For example, it would be meaningful to vary all second-order transport coefficients of the hydrodynamic model. In this work we took only a small step in this direction by varying the shear relaxation time.

Other model uncertainties/biases are more difficult to study. For example, our ansatz of initial conditions for hydrodynamics may not be sufficiently flexible to capture all relevant features of the early pre-hydrodynamic evolution in heavy ion collisions. Accounting for these model limitations systematically is still challenging.

In Ref. [27] a parametrized systematic model discrepancy was included; this single percentage uncertain parameter was meant as a proxy for all systematic model discrepancies. This parameter was added in quadrature to the covariance matrix of the Gaussian process for each principal component, in the form of a diagonal matrix parametrized by the single parameter σ_m . That is, to every principal component was added the same systematic uncertainty in percentage. This results in a complicated distribution of the uncertainty across observables, depending on the linear transformation from principal components to observables.

While the above approach may be a step in the right direction, we worry about the crudeness of this distribution of systematic uncertainty across the observables. In consequence, we did not include this type of uncertainty in our current analysis.

D. Sampling of the posterior

In our case the posterior is a 18-dimensional probability distribution.¹⁹ Estimation of the posterior is accomplished via Markov Chain Monte Carlo algorithms [132]. Typical algorithms, including the Metropolis-Hastings algorithm, are able to estimate the shape of the posterior without knowledge of its normalization.

Efficient and accurate Markov Chain Monte Carlo algorithms are now readily available, thanks to their widespread use in other fields (e.g. in cosmology). This includes nested sampling, Hamiltonian methods, and parallel tempering [133]. In this work, we used an implementation of parallel tempering [134]; the algorithm showed good convergence

in sampling our posterior and at the same time made it possible to estimate the Bayesian evidence. The latter is discussed in Section XI. While this algorithm does provide an estimate for the evidence we will use this information only for performing model selection; for parameter estimation the unnormalized posteriors were used and plotted.

E. Maximizing the posterior

Although the primary result of parameter estimation is the posterior distribution, it is also useful to calculate the point in parameter space which maximizes the posterior. This is referred to as the Maximum a Posteriori (MAP) set of parameters. Because throughout this work we use priors which are uniform distributions, the MAP parameters are those which maximize the likelihood function; that is, the parameters which optimize the fit to the experimental data.

VI. CLOSURE TESTS

Closure tests are required to ensure that, in a situation with known model parameters, the Bayesian inference framework correctly reproduces them from a set of “mock data”. These data are obtained by running the model to generate predictions instead of real data for the observables which one intends to use for the model calibration. Ideally, the posterior of a closure test should approach a delta-function around the true value of the model parameters, $\mathcal{P}(\mathbf{x}|\mathbf{y}_{\hat{\mathbf{x}}}) \rightarrow \delta(\mathbf{x}-\hat{\mathbf{x}})$. In practice, this delta function is always smeared by the uncertainties present in the Bayesian parameter estimation.

A first source of uncertainties is in the model calculations used instead of experimental data: since for heavy-ion collisions the initial conditions fluctuate stochastically and running the model is expensive, statistical uncertainties of the “mock data” are often larger than those from real experiments, and these will propagate non-trivially and contribute to the width for the parameter posterior. Additional uncertainties are contributed by the emulator: (i) statistical uncertainties from the calculations used to train the emulator; (ii) interpolation uncertainty from the limited number of parameter samples used to train the emulator; and (iii) the limited number of principal components that are interpolated via Gaussian processes. Finally, there might be partial degeneracies in the model which makes it difficult to match a unique set of observables to a unique set of model parameters. Even if a sufficiently large set of observables is used to avoid exact degeneracies, approximate degeneracies can persist until all the uncertainties decrease below a certain threshold.

Closure tests provide a way to identify such potential issues and, for a chosen set of observables, quantify the effect of these types of uncertainties on the parameter estimation before any comparison with measurements is performed. Closure tests can also help clarify the level of constraint on the model parameters that can be expected given the emulator uncertainties. These two aspects of closure tests are not independent; however they are sufficiently different objectives that they benefit

¹⁹ There are 16 shared parameters and one additional parameter per collision system (the T_RENTO normalization).

being discussed separately.

A. Validating Bayesian inference with closure tests

In what follows we show a sample set of closure tests. They employ the same emulator as will later be used for comparison with experimental data.

We proceed as follows:

1. We generate a set of design points $(\mathbf{x}_i; i = 1, \dots, m_v)$ for training the model emulator, and a separate set of design points for validation $(\hat{\mathbf{x}}_i; i = 1, \dots, \hat{m}_v)$.
2. We perform full model calculations at both the training and validation design points and compute final state observables.
3. We perform principal component analysis on the *training* calculations, and fit a Gaussian process to each retained principal component.
4. For each point i in the *validation* set, we use the trained emulator to perform parameter estimation using the calculated model observables at validation point $\hat{\mathbf{x}}_i$ as the “data”.
5. We compare the posterior $\mathcal{P}(\mathbf{x}|\mathbf{y}_{\hat{\mathbf{x}}_i})$ to the known true values $\hat{\mathbf{x}}_i$.

Our emulator uses 500 design points. At each design point we use the full model to compute predicted values for all observables that will also be used in the calibration with real data (see Sec. VII). As discussed previously, our model includes statistical fluctuations, which arise from averaging over a finite number of initial conditions (2500 hydrodynamic events per design point), as well as Cooper-Frye sampling each particleization hypersurface a finite number of times (at least 10^5 particles sampled per hydrodynamic event). We use 10 principal components, which explain approximately 98% of the model variance for Pb-Pb data at $\sqrt{s_{\text{NN}}} = 2.76$ TeV. These uncertainties, combined with the emulator uncertainty discussed above, lead to a finite spread of our posterior $\mathcal{P}(\mathbf{x}|\mathbf{y}_{\hat{\mathbf{x}}_i})$. What we can verify is how often the true parameters lie within given regions of credibility.

Figure 3 shows the result of our closure tests for 9 sets of validation points. We focus on the specific shear and bulk viscosities of the QGP, η/s and ζ/s . The parametrization of these physical quantities involves non-linearly correlated parameters. These parameters are no more than a few degrees of freedom we put into the functional form of $(\eta/s)(T)$, $(\zeta/s)(T)$; therefore, it is of less physical importance to focus on the posterior of these parameters individually. Instead, what we show in Fig. 3 are their resultant posterior for η/s and ζ/s as functions of T , compared to the assumed “true values” shown as dashed black lines). Red and blue bands show the 60% and 90% credible intervals of the estimation; at different temperatures these credible intervals are determined independently. The results demonstrate that the functional shapes of the “true” viscosity-to-entropy ratios are well enclosed by the inferred 60% and 90% confidence regions.

B. Guiding analyses with closure tests

Figure 3 provides convincing evidence that the emulators and the Bayesian analysis are performing well. Importantly, it also provides a wealth of information on the eventual results of the Bayesian parameter estimation when performed with real data.

Recall that the result of the Bayesian analysis depends on a variety of factors, including (i) the choice of observables, (ii) the uncertainty on these observables, and (iii) the uncertainty of the emulator. In an ideal world, the emulator uncertainty would be much smaller than the uncertainties on the observables; in such a scenario, the result of the Bayesian analysis would be essentially independent of the emulator. This is always the goal one should strive for; however, in practice this is difficult to achieve, and most Bayesian parameter estimations are performed under much less ideal conditions.

For the case shown in Fig. 3, emulator uncertainties are not negligible. However, given the emulator we use in this work, a comparison of the closure test in Fig. 3 with the prior from Fig. 2 demonstrates that the current Bayesian parameter estimation method, as well as the selection of observables, have the best constraining power for η/s and ζ/s at low temperatures. This is expected since these temperatures are closer to the switching temperature between hydrodynamics and the hadronic transport model, and much of the space-time volume explored by the expanding medium is characterized by such moderate temperatures [135]. The closure tests also indicate that the uncertainty on the viscosities is large at higher temperatures. We believe this could be a consequence of the smaller amount of time spent by the systems at high temperatures [135], decreasing the sensitivity of the observables’ response to the transport parameters in this temperature region.

Additional observables or collision energies may help improve these constraints on the viscosities of QCD. For example, emission of electromagnetic radiation puts somewhat stronger weight on the earlier and shorter-lived hot fireball regions than hadrons do [135]. On the other hand, electromagnetic observables are plagued by larger statistical and systematic uncertainties. Closure tests can be used exactly for the purpose of assessing the value of adding such additional measurements even before such data are available: they allow for quantifying the contribution of different observables towards constraining the properties of the quark-gluon plasma. In the future this could be an important tool to guide the priorities of experimental campaigns. Observables contribute differently to constraining different model parameters: by quantifying the effect of adding a new observable, or reducing the uncertainty on an existing one, one can provide meaningful feedback which measurements should be prioritized. These methods are closely related to those employed in “Bayesian Experimental Design” [136].

One caveat to be kept in mind in this context is that closure tests evidently rely on the correctness of the underlying physics model. When we compare to actually experimentally observed data, we cannot assume that our model provides an exact description of the observables even when given the right choice of parameters [131]. The systematic model discrepancy must




FIG. 3. Closure tests of the specific shear (left) and bulk (right) viscosities using 9 sets of validation points in the parameter space. Performed with the emulator for Pb-Pb $\sqrt{s_{\text{NN}}} = 2.76$ TeV collisions. Shown in blue are the 90% credible intervals and in red the 60% credible intervals.

not be forgotten. Hence, the result of a closure test should not be taken as the final word: the importance of a given observable in constraining model parameters may need to be revisited when physics tested by this observable is modified in the model.

In spite of these unavoidable limitations, closure tests can provide important guidance to experimental collaborations to help determine which observables can best constrain physical parameters.

VII. BAYESIAN PARAMETER ESTIMATION USING RHIC AND LHC MEASUREMENTS

In this Section we perform a Bayesian parameter estimation with RHIC and LHC measurements. We focus on constraints for the shear and bulk viscosities provided by transverse-momentum-integrated data from LHC and RHIC. We perform these first analyses for a specific model of viscous corrections at particlization, the Grad model (see Section III C). The effect of using different viscous corrections as well as other systematic uncertainties of the model are quantified in the next section.

A. Constraints on η/s and ζ/s from Pb-Pb measurements at $\sqrt{s_{\text{NN}}} = 2.76$ TeV

We first study the parameter estimates including only the data from Pb-Pb collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV. We use the following measurements from the ALICE collaboration:

- the charged particle multiplicity $dN_{\text{ch}}/d\eta$ [137] for bins in 0–70% centrality;

- the transverse energy $dE_T/d\eta$ [138] for bins in 0–70% centrality;
- the multiplicity dN/dy and mean transverse momenta $\langle p_T \rangle$ of pions, kaons and protons [139] for bins in 0–70% centrality;
- the two-particle cumulant harmonic flows $v_n\{2\}$ for $n = 2, 3, 4$, for bins in 0–70% centrality for $n = 2$, and for bins in 0–50% centrality for $n = 3$ and 4 [140];
- the fluctuation in the mean transverse momentum $\delta p_T/p_T$ [141] for bins 0–70% centrality.

Before being reduced by principal component analysis this data set represents 123 “independent observables”, given that measurements at different centralities are treated as separate observables. We found that 10 principal components (linear combinations of observables) are sufficient to capture most of the sensitivity of these observables to the full set of parameters: they capture more than 98% of the variance. This number of dominant principal components represents 8% of the total number of observables. Thus there is a significant amount of redundant information in the observables with respect to our model parameters. We highlight that we tested the effect on our analysis of reducing the number of principal components: we determined that our results are robust with respect to the number of principal components used. The results of this test are presented in Appendix C.

We note that all observables used in this analysis are p_T -integrated. Observables which are differential in transverse momentum undeniably carry additional information about the medium [30, 31]. There is reasonable evidence that low- p_T ($p_T \lesssim 1.5$ GeV) information is generally included in p_T -integrated observables [23]. The higher- p_T range ($p_T \gtrsim$

1.5 GeV) tends to have larger modeling uncertainties, if only from viscous corrections at particlization which can be very significant at higher transverse momenta. At sufficiently high p_T , hadron production is beyond the realm of hydrodynamics altogether; this threshold is not known precisely, but even a breakdown at $p_T \gtrsim 2 - 3$ GeV would not be wholly surprising. Because of these limitations, there is a risk that inferences using observables in the higher- p_T range ($p_T \gtrsim 1.5$ GeV) lead to more precise but unreliable constraints on the parameters. While both avenues are worth exploring, in the present analysis we opt for the more conservative approach of using p_T -integrated observables that introduce less model bias, while also studying in greater detail model uncertainties.

The posteriors for the shear and bulk viscosities are shown in Fig. 4. Recall that this result is for a *single* viscous correction model, the Grad viscous correction.

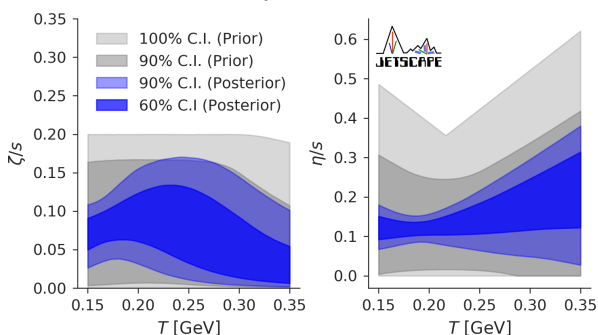


FIG. 4. The posterior for specific bulk (left) and shear (right) viscosities resulting from a Grad viscous correction model parameter estimation using ALICE data for Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV

We first note a general feature which will remain when we examine other viscous corrections and include more systems: the constraint on the shear and bulk viscosities is best near the switching temperature T_{sw} . This was already observed in the closure tests performed in Sec. VI. The viscous corrections in the particlization procedure depend on the magnitude of shear stress $\pi^{\mu\nu}$ and bulk pressure Π on the switching surface, making the model predictions sensitive to the viscosities near these temperatures. As we have discussed in the closure test, the uncertainties in ζ/s and η/s are larger in the high temperature region. We see that for the bulk viscosity in particular, our 90% posterior credible interval is only slightly smaller than our prior above 250 MeV.

B. Constraints on η/s and ζ/s from Au-Au measurements at $\sqrt{s_{NN}} = 0.2$ TeV

We also examine the constraints on the viscosities provided by the existing data for Au-Au collisions at $\sqrt{s_{NN}} = 200$ GeV. Heavy-ion collisions at RHIC provide complimentary information, having smaller temperatures and a shorter lifetime than collisions at the LHC. We use the following experimental measurements from the STAR Collaboration:

- the yields dN/dy and mean transverse momenta $\langle p_T \rangle$ of pions and kaons for bins in 0–50% centrality [142];
- the two-particle cumulant harmonic flows $v_n\{2\}$ for $n = 2, 3$ for bins in 0–50% centrality [143, 144].

We remark that because of the tension between STAR and PHENIX measured proton yields at mid-rapidity in Au-Au collisions at $\sqrt{s_{NN}} = 200$ GeV [142, 145], we have deliberately left the proton yield and mean transverse momentum out of the current comparison.²⁰

The above includes 29 observables, again counting centralities as separate observables. After performing principal component analysis, we kept 6 principal components (equivalent to 21% of the total number of observables), which explain more than 98% of the variance of the observables across the parameter space.

The estimated viscosities using only these measurements from RHIC, again for the Grad viscous correction, are shown in Fig. 5. The posteriors for specific bulk and shear viscosity

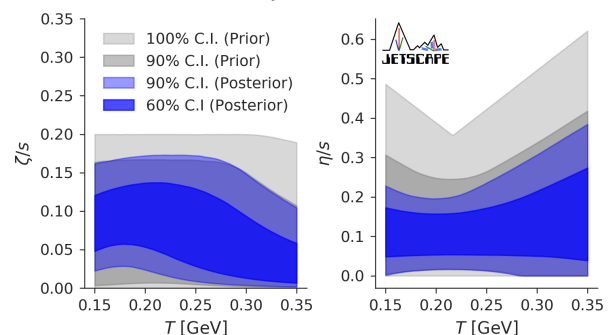


FIG. 5. The posterior for specific bulk (left) and shear (right) viscosities resulting from a model parameter estimation using STAR data for Au-Au collisions at $\sqrt{s_{NN}} = 200$ GeV.

when calibrating against only RHIC data have in general different features than those given by the LHC data. For instance, we see that a large specific bulk viscosity is allowed near the switching temperature. Also, the 90% credible interval for the specific shear viscosity extends to lower values for these data than the LHC data; only using these RHIC observables, a specific shear viscosity which is nearly zero ($\eta/s < 0.03$) is consistent with the data.

It is important to note that not only the specific bulk and shear viscosity parameters have different posteriors, but in general the entire parameter posterior will be different when we use RHIC observables rather than LHC observables. The two are compared for a different subset of model parameters in Appendix B.

²⁰ Moreover, both measurements [142, 145] show notable excess of proton production over anti-proton production, suggesting the importance of including a non-zero baryon chemical potential (μ_B) in our calculation. The current study assumes $\mu_B = 0$ in both initial condition and dynamical evolution, and improvement should be considered in future studies.

C. Viscosity estimation and model accuracy for combined RHIC & LHC data

Reviewing Figs. 4 and 5 we find that the observables at the LHC give stronger constraints on the slope of the specific shear viscosity at large temperature. It is the general expectation that higher $\sqrt{s_{\text{NN}}}$ collisions at the LHC are more sensitive to the transport coefficient at high temperature. This conclusion was verified quantitatively in previous Bayesian parameter estimation [24, 146]. For the present analysis, we do caution that we currently use a different number of observables at RHIC and the LHC; consequently, we are not in a position to compare systematically the constraining power of the two collision energies at the moment. We do expect RHIC and LHC data to be complementary, and we proceed to a combined Bayesian parameter estimation for Pb-Pb at $\sqrt{s_{\text{NN}}} = 2.76$ TeV and Au-Au at $\sqrt{s_{\text{NN}}} = 200$ GeV collisions. For this combined analysis, the viscosity posterior for the Grad viscous correction is shown in Fig. 6.




FIG. 6. The posterior for specific bulk (left) and shear (right) viscosities resulting from a model parameter estimation using combined data for Au-Au collisions at $\sqrt{s_{\text{NN}}} = 200$ GeV and Pb-Pb collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV.

As discussed in Section V A, all parameters are held the same for the two systems except for their overall normalizations of the initial conditions — $N[2.76 \text{ TeV}]$ and $N[0.2 \text{ TeV}]$. Recall that model parameters being kept constant does not imply that the effective physical quantities are the same at RHIC and the LHC. For example, the transport coefficients are temperature dependent, and the free-streaming time depends on $\sqrt{s_{\text{NN}}}$ and centrality through the total energy of the event.

The information gained by fitting both systems slightly reduces the width of the credible intervals for the specific shear and bulk viscosities at temperatures above 250 MeV; the 90% confidence band in the posterior for specific shear and bulk viscosity is slightly smaller than the credible intervals given by calibrating against either one of these two systems alone. This illustrates the added constraining power accessed by combining the two data sets.

The simultaneous fit to experimental observables is shown in Fig. 7, where we have plotted the emulator prediction for the observables at one hundred parameter samples drawn randomly from the posterior. Note that, in spite of some undeniable tension in the simultaneous fit of ALICE and STAR data

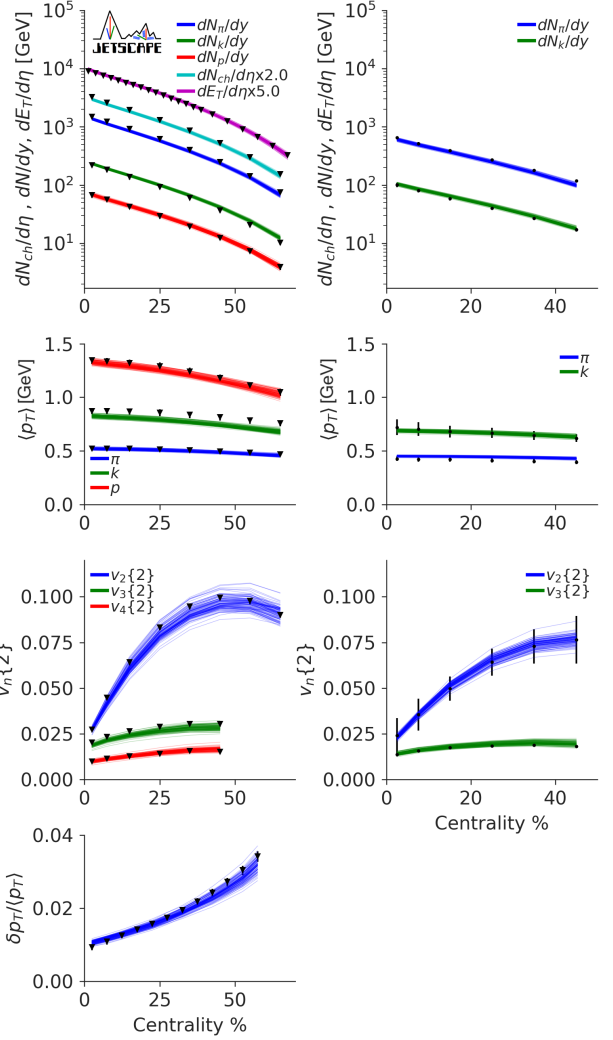


FIG. 7. The observables predicted by the Grad viscous correction emulator, drawn from the posterior resulting from the combined fit of ALICE data (left) for Pb-Pb collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV and STAR data (right) for Au-Au collisions at $\sqrt{s_{\text{NN}}} = 200$ GeV. The simultaneous fit yields model observables which agree within $\sim 20\%$ of experimental measurements.

(for example in the mean transverse momenta of kaons), our hybrid model can describe simultaneously all of the observables we considered for the two systems to within 20% of the experimental results. As discussed earlier, this is important: our confidence in the significance of this section's parameter estimates rests on a good description of the experimental data when sampling model parameters according to their posterior probability distribution.

As a final emulator validation, we have calculated the Maximum A Posteriori (MAP) parameters of the Grad viscous correction model. Using these parameters, we simulated 5,000 fluctuating events and performed centrality averaging. The comparison between the hybrid model prediction at the MAP parameters and the experimental data are shown in Fig. 8, and MAP parameters for the Grad, Chapman-Enskog and Pratt-

Torrieri-Bernhard models are listed in Table II.²¹

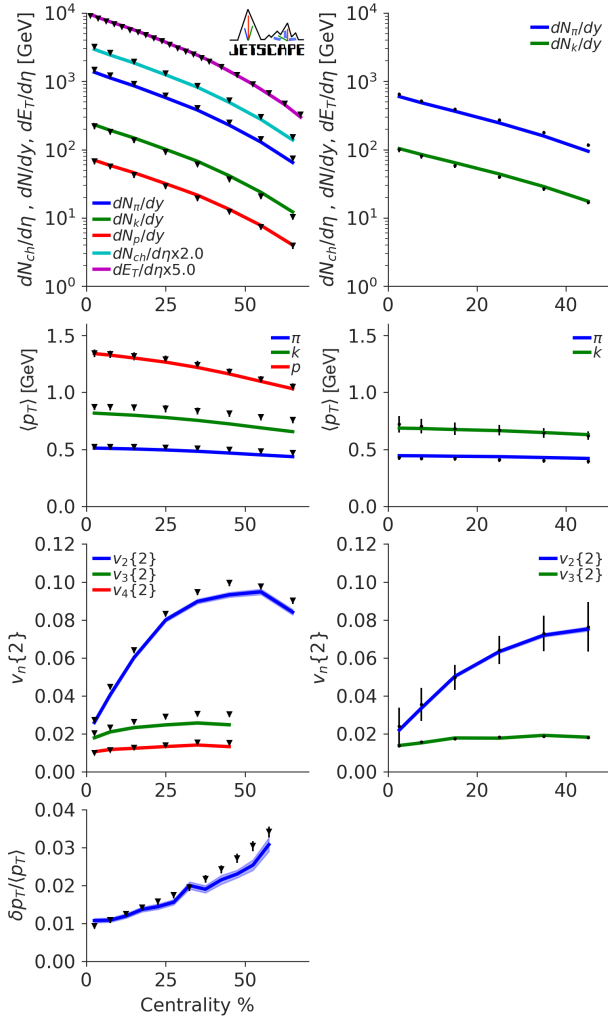


FIG. 8. The observables resulting from averaging over five thousand fluctuating events for each system, run with the MAP parameters of the combined calibration of ALICE data for Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV and STAR data for Au-Au collisions at $\sqrt{s_{NN}} = 200$ GeV. Results are shown for the Grad viscous correction. Shaded bands around model predictions reflect the variance arising from initial state fluctuations combined with statistical fluctuations from the particlization sampler. Pb-Pb $\sqrt{s_{NN}} = 2.76$ TeV events are shown at left and Au-Au $\sqrt{s_{NN}} = 0.2$ TeV events at right.

Because our prior for each of these parameters was uniform on a finite range, the parameters which maximize the posterior also maximize the likelihood function; this means that they also optimize the fit to the experimental data (i.e. minimize χ^2).

²¹ For reasons explained in Sec. VIII A, the Pratt-Torrieri-McNelis model (PTM) is left out from Table II.

Parameter	Grad	CE	PTB
$N[2.76 \text{ TeV}]$	14.2	15.6	13.2
$N[0.2 \text{ TeV}]$	5.73	6.24	5.31
p	0.063	0.063	0.139
σ_k	1.05	1.00	0.98
w [fm]	1.12	1.19	0.81
d_{\min}^3 [fm ³]	2.97	2.60	3.11
τ_R [fm/c]	1.46	1.04	1.46
α	0.031	0.024	0.017
T_η [GeV]	0.223	0.268	0.194
a_{low} [GeV ⁻¹]	-0.776	-0.729	-0.467
a_{high} [GeV ⁻¹]	0.37	0.38	1.62
$(\eta/s)_{\text{kink}}$	0.096	0.042	0.105
$(\zeta/s)_{\text{max}}$	0.133	0.127	0.165
T_ζ [GeV]	0.12	0.12	0.194
w_ζ [GeV]	0.072	0.025	0.026
λ_ζ	-0.122	0.095	-0.072
b_π	4.65	5.62	5.54
T_{sw} [GeV]	0.136	0.146	0.147

TABLE II. Table of MAP parameters of the Grad, Chapman-Enskog (CE) and Pratt-Torrieri-Bernhard (PTB) viscous correction models, from combined RHIC and LHC data.

VIII. PARAMETER ESTIMATION AND SYSTEMATIC MODEL UNCERTAINTIES

In this section, we continue our exploration of the estimated parameter posterior for the combined LHC Pb-Pb $\sqrt{s_{NN}} = 2.76$ TeV and RHIC Au-Au $\sqrt{s_{NN}} = 0.2$ TeV data. We identify and discuss some of the largest sources of theoretical uncertainty in the physical model, and the effect these uncertainties have on constraining the viscosities of QCD.

The first source of uncertainty that we investigate in Section VIII A originates from mapping the hydrodynamic fields to hadronic momentum distributions, the “viscous corrections” at particlization, discussed in Section III C. Recall that the results from the previous section were for a specific choice of viscous corrections, the Grad model.

The viscous corrections correspond to one source of uncertainty in the transition from hydrodynamics to particles. A second source is the determination of the particlization hypersurface, which in this work is defined at a fixed switching temperature T_{sw} . We discuss the dependence of our results on this switching temperature in Section VIII B. We discuss at the same time the transition between the early stage of the model and hydrodynamics, which we find exhibits clear correlation with the switching temperature.

Finally, we discuss the effect of second-order transport coefficients, as quantified with the shear relaxation time in Section VIII C.

A. Mapping hydrodynamic fields to hadronic momentum distributions

As discussed in Section III C, there are still significant uncertainties in matching the energy-momentum tensor from hydrodynamics to a unique hadronic momentum distribution. In-




FIG. 9. The 90 % credibility intervals for the prior (gray shaded area) and for the posteriors (colored outlines) of the specific bulk (left) and shear (right) viscosities, for three viscous correction models: Grad (blue), Chapman-Enskog (CE, red) and Pratt-Torrieri-Bernhard (PTB, green). The Pratt-Torrieri-McNelis (PTM) posterior is not shown, but is nearly identical with the Chapman-Enskog result.

evitably this affects the results of the phenomenological constraints we obtain on the shear and bulk viscosity of QCD: the posterior for every model parameter depends on the choice of viscous correction at particlization. Recall that in this work, we chose to study four different models of viscous corrections (see Section III C): (i) Grad (“14-moments”); (ii) Chapman-Enskog in Relaxation Time Approximation; (iii) an exponentiated version of the Chapman-Enskog model referred to as “Pratt-Torrieri-McNelis”; and (iv) an additional exponentiated model of viscous corrections referred to as “Pratt-Torrieri-Bernhard”. In our tests, we found that the posteriors for the exponentiated Chapman-Enskog ansatz called “Pratt-Torrieri-McNelis” were always very similar to the results for the linearized Chapman-Enskog ansatz. To avoid clutter we therefore decided not to show the posteriors for the Pratt-Torrieri-McNelis model, neither in this Section nor anywhere else in this work. We begin with the marginalized posteriors for the QGP viscosities, shown in Fig. 9. The Figure exhibits clear differences in the experimentally preferred shear and bulk viscosities for the different viscous correction models. Remember that it is important to read Fig. 9 with respect to the parameter prior whose 90% credibility region is indicated by the gray shaded area. A posterior that covers the same area as the prior should be interpreted as indication for weak or even absence of experimental constraints. On the other hand, a posterior that systematically excludes certain regions of the prior provides good evidence that parameter values in these excluded regions are disfavored by data.

For the bulk viscosity (left), we can see in Fig. 9 that each of the different viscous correction models excludes only relatively small regions of the prior. For all four particlization models the constraints on ζ/s are tighter at lower temperatures than at higher ones. However, the ζ/s regions favored by each model at low temperature differ from each other: the Grad viscous correction model favors a larger ζ/s where the Chapman-Enskog model favors lower values, with the Pratt-Torrieri-Bernhard model lying in between. We note in particular that the Pratt-Torrieri-Bernhard posterior is very narrow at

low temperature. We understand this to be a consequence of mean transverse momenta and harmonic flows being very sensitive to the specific bulk viscosity near the switching temperature for the Pratt-Torrieri-Bernhard viscous correction model. We quantify and revisit this difference in sensitivity of the viscous correction models in Section X.

Overall, only large values of ζ/s at low temperature are excluded by all three viscous correction models. As such, our constraints on the bulk viscosity are limited, especially after accounting for the model uncertainty introduced by the viscous corrections.

For the shear viscosity shown in the right panel of Fig. 9 we encounter a similar situation: limited constraints on η/s at higher temperatures, and exclusion of large values of η/s at low temperature by all viscous correction models. Overall, shear viscosity is best constrained at temperatures around 200 MeV.

From the results of this section, we see that viscous corrections represent a considerable uncertainty in constraining the QGP shear and bulk viscosities. It is important to remember that all viscous correction models studied in this work are based on relatively simple assumptions. The capacity of any of these models to describe correctly the momenta and chemistry of a realistic out-of-equilibrium system of hadrons is still under investigation (see Ref. [97, 98] and references therein for a recent overview). For instance, all of these particlization models assume that the hydrodynamic shear stress is shared “democratically” among the hadronic species. This approximation greatly simplifies the models, but microscopic transport theory suggests that it may not be suitable for heavier hadrons such as protons [96]. Additional theoretical efforts (see e.g. Refs. [95–98, 147]) may be able to shed more light on this question and provide additional insights that can be used for tightening our prior assumptions in future Bayesian analyses. Until this happens the particlization model uncertainty must be considered as “irreducible” and is best accounted for by Bayesian Model Averaging as reported in [28].

B. Transition to and from hydrodynamics: initial state and switching temperature

The previous section focused on the uncertainty originating from transitioning from hydrodynamics to a particle description of the system. This transition occurs on a hypersurface defined by a temperature T_{sw} . Recall that this switching temperature is also a model parameter, allowed to vary between 135 and 165 MeV.

The other transition point to hydrodynamics is the time at which hydrodynamics is initialized with the energy-momentum tensor from the preceding free-streaming evolution (Section III A). This hypersurface is defined at a constant proper time τ_{is} , the value of which depends on two parameters as defined in Eq. (19). The hydrodynamic initial conditions on this hypersurface further depend on the initial condition parameters of the T_RENTO ansatz. In this section, we discuss the posterior of T_{sw} , τ_{is} and the T_RENTO parameters, how they are correlated, and how they are affected by the viscous correction

models discussed in the previous section.

Fig. 10 provides a dimensionally reduced representation of the joint posterior probability distribution for all model parameters except those related to shear and bulk stress, for two viscous correction models, Grad (blue) and Chapman-Enskog (red). Histograms on the diagonals are the marginalized one-dimensional posteriors for each parameter. The off-diagonal histograms are the joint posteriors for each pair of parameters, marginalized over all others.

One observes that, within the chosen prior range, the normalizations of the initial energy density for the two systems $N[2.76 \text{ TeV}]$ and $N[0.2 \text{ TeV}]$ are well constrained by the observables, for both viscous correction models, but with slightly shifted peak values. Note that, since the final multiplicities are fixed by experiments, lower normalization factors for the initial energy (and hence entropy) density reflect larger viscous heating effects during the subsequent dynamical evolution. The amount of viscous heating is also affected by the particlization temperature T_{sw} , with lower values of T_{sw} corresponding to longer lifetimes of the hydrodynamic stage.

We further find that the estimation of the generalized mean parameter p is the same for the two viscous correction models, close to $p = 0$. We verified that p is also close to zero ($p \approx 0.1$) with the Pratt-Torrieri-Bernhard viscous model. These p -values are consistent with previous studies which also used the Pratt-Torrieri-Bernhard viscous correction model but differed in some other model details [53]. The result $p \approx 0$ seems to be remarkably robust across all existing Bayesian inference analyses of high-energy heavy-ion collision data [26, 27, 53, 67]. We note that T_{RENT}o with $p = 0$ shares important aspects of fluctuating collision geometry with phenomenologically successful initial condition models based on saturation physics. For example, $p = 0$ predicts that the energy deposition is proportional to $\sqrt{T_A T_B}$ as discussed in Sec. III A. This unique feature of the initial local energy density being solely a function of the product $T_A(x_\perp)T_B(x_\perp)$ is also found in the pQCD+saturation based EKRT initial condition model [81],²² and in models with approximate longitudinal boost-invariance the $\sqrt{T_A T_B}$ dependence can be motivated by simple arguments based on conservation of energy and momentum during the initial energy deposition process [69]. Earlier studies [26, 44] further noted that for $p \approx 0$ T_{RENT}o can reproduce the centrality dependent 2-particle cumulant eccentricity ϵ_2 and triangularity ϵ_3 of the IP-Glasma initial condition model [68]. However, one should keep in mind that the two models have very different participant scaling of local energy deposition. According to Eqs. (2,3), T_{RENT}o for $p = 0$ sets the initial local energy density proportional to $\sqrt{T_A T_B}$, but the IP-Glasma model predicts a $T_A T_B$ scaling immediately after the collision [148]. The two models also have different levels of granularity and fluctuation in the energy deposition [68]. Moreover, studies [149, 150] that used the IP-Glasma model

to initialize the hydrodynamics defined centrality differently from the present and earlier studies using the T_{RENT}o model [26, 67]. All these differences convoluted into the comparison of centrality dependent ϵ_2 and ϵ_3 between the two models. Therefore, T_{RENT}o ($p = 0$) should not be considered a substitute for these theories based on saturation physics but rather taken as an efficient parametrization of general geometric features shared by these initial state models that happen to be preferred by the experimental data.

The nucleon width w , which controls the transverse length scale of energy fluctuations in the initial state, is also well constrained by the data and found to be about 1.1 fm, nearly independent of the two viscous correction models. A smaller value for this nucleon width ($w \approx 0.9 \text{ fm}$) was found, however, with the Pratt-Torrieri-Bernhard viscous correction.

In general, our conclusions for the T_{RENT}o parameters is that they do not appear to be highly sensitive to the choice of viscous correction model at particlization. However, the particlization uncertainty should not be entirely ignored as it can be larger than the width of the posteriors for each of these parameters.

The posteriors for the free-streaming time scale τ_R and the associated energy dependence parameter α are not easily interpreted; they are correlated by our parametrization (Eq. (19)) of the effective free-streaming time τ_{fs} . We can point out, however, that the posteriors for the Grad and Chapman-Enskog models are quite different for these two parameters; in the case of the Chapman-Enskog model, the posterior for α is bimodal. It is not clear whether the peak in α near -0.3 is a local maximum or if there exists a global maximum in the posterior for values of α less than -0.3 . We cannot currently differentiate between these two scenarios.

We can take a closer look at the posterior of the free-streaming time by plotting it as a function of the physical scale e_0 , which is the magnitude of the average initial energy density in the transverse plane. This is shown in Fig. 11. We see that the 90% credible interval for the energy dependence of the free-streaming time is not well constrained, and that it is consistent with having no energy dependence. What is constrained is the overall magnitude of the free-streaming time. The Chapman-Enskog model has a posterior which prefers smaller free-streaming times, while the Pratt-Torrieri-Bernhard model prefers the largest free-streaming time of all particlization models studied.

It is generally expected that collisions with higher energy density will hydrodynamize more rapidly [72]. In our model this would correspond to $\alpha < 0$. The peak at $\alpha > 0$ in the posterior for α is at variance with this expectation. One should remember, however, that our pre-hydrodynamic model does not naturally lead to hydrodynamization which is instead enforced by hand at τ_{fs} . As such, it is conceptually problematic to associate our free-streaming time τ_{fs} with a hydrodynamization time. As discussed in connection with Eq. (18), hard-matching an energy-momentum tensor from a conformally invariant pre-hydrodynamic evolution model without thermalization to dissipative hydrodynamics with a non-conformal EoS leads to a (possibly large) *positive* (i.e. unphysical) initial value for the bulk viscous pressure whose subsequent decay can have

²² Though the EKRT model used a different parametrization for the relation between energy density and $T_A T_B$, its functional form agrees very well with the $\sqrt{T_A T_B}$ relation for typical nuclear thickness functions obtained for lead nuclei.

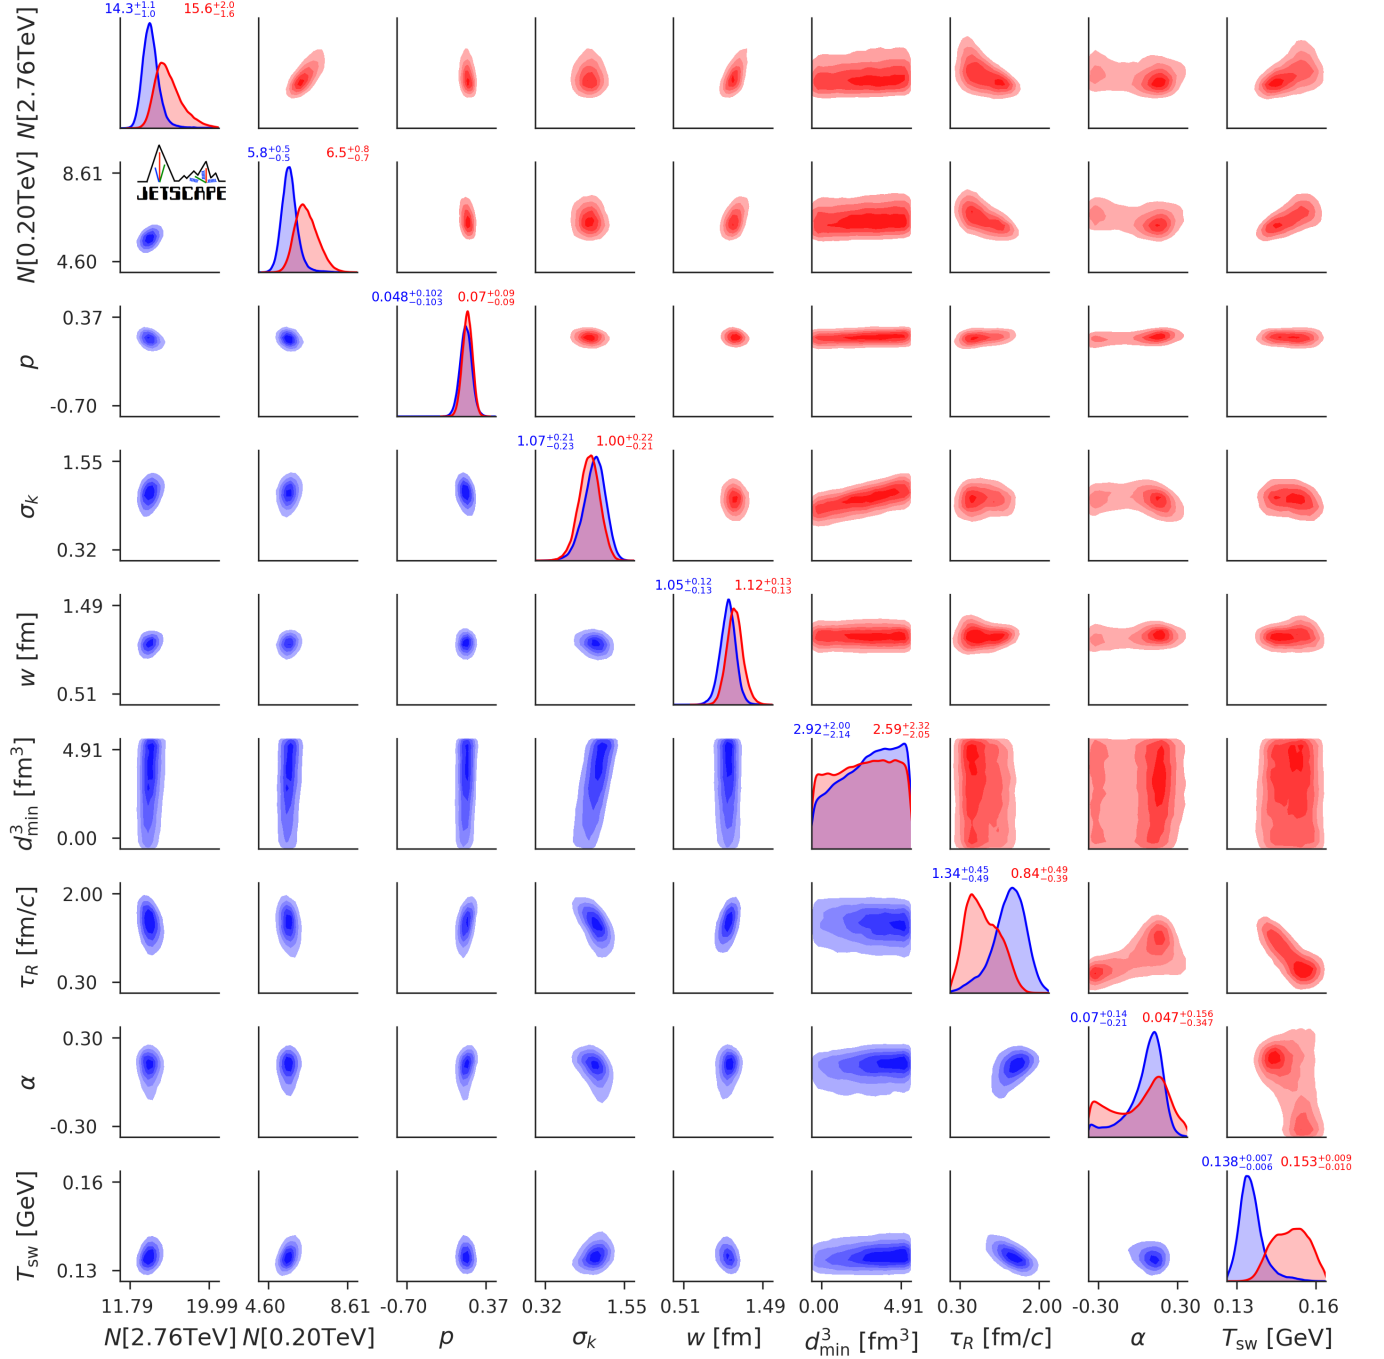


FIG. 10. The posterior for Grad (blue) and Chapman-Enskog (red) viscous corrections for select parameters related to the initial state, pre-hydrodynamic evolution and switching temperature. The histograms on the diagonal are the marginal distributions for each parameter, with appended numbers denoting the median and the left and right limits of the 90% credible interval. Off-diagonal histograms display the joint-posterior of each pair of parameters, marginalized over all others.

counter-intuitive effects on the subsequent hydrodynamic flow and its dependence on τ_{fs} . Recent studies demonstrate that this problem persists when the free-streaming module is replaced by a thermalizing but conformal effective kinetic theory, and that the magnitude of the mismatch depends on centrality [151]. Although we have not been able to fully dissect

the mechanisms leading to positive preferred values for α in our analysis, we strongly suspect that these issues play a role.

Turning to the later stages of the collision, we now look at the posterior for the switching temperature in Fig. 10. Its marginalized posterior turns out to be quite different for the two particlization models. We find that for the selected exper-

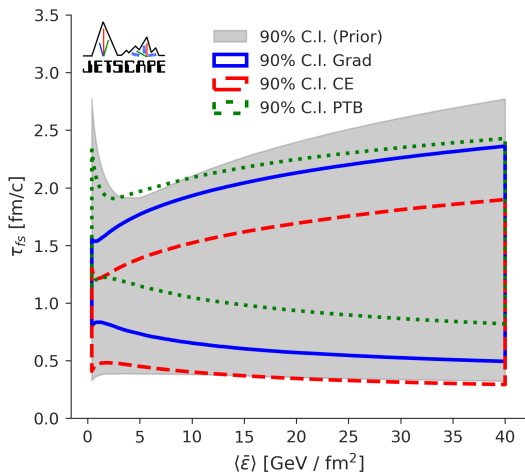


FIG. 11. The 90% posterior credible intervals for the free-streaming time, as a function of the initial average transverse energy density defined in Eq. (20), resulting from parameter estimation using combined data for Au-Au collisions at $\sqrt{s_{NN}} = 200$ GeV and Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV. The Grad model is shown with solid blue, Chapman-Enskog with dashed red, and Pratt-Torrieri-Bernhard with dotted green lines.

imental observables the effects of increasing the magnitude of the bulk viscous pressure or increasing T_{sw} are qualitatively similar. We verified that if we hold all other parameters fixed while increasing the switching temperature from 135 MeV to 165 MeV, the mean transverse momenta of pions and protons is reduced and the number of protons is increased. On the other hand, holding T_{sw} fixed and increasing $(\zeta/s)_{max}$ has the same effect. Because the Grad model prefers a large specific bulk viscosity near switching, it also prefers a lower switching temperature.

Although some of the parameters which define the $T_{RENT}o$ model are well constrained, their interpretation is not always straightforward. To what extent $p \approx 0$ in the $T_{RENT}o$ model provides support for saturation physics or is mostly a consequence of energy-momentum conservation combined with approximate boost invariance at high collision energies deserves further study. Similarly, we found our posterior for the energy density dependence of the free-streaming time difficult to understand. Further theoretical understanding of these issues is likely to find its way into improved models. For example, there is considerable room for improving the description of the pre-hydrodynamic evolution stage. It is encouraging that, using the likely values for the $T_{RENT}o$ model parameters, we are able to describe our experimental observables with good accuracy. Still, there is obvious value in seeking models which can not only fit the experimental data, but at the same time offer a coherently and consistently interpretable physical picture.

C. Second-order transport coefficients: shear relaxation time

As discussed in Section III B, second-order transport coefficients are treated differently in this work than first-order




FIG. 12. The posterior of specific bulk(left) and shear(right) viscosities, depending on whether one marginalizes over the shear relaxation time factor b_π (solid blue) or fixes it (dashed or dotted blue). The shear relaxation time and magnitude of η/s are seen to be inversely related when fitting the LHC and RHIC data.

ones: the shear and bulk viscosities are parametrized, while the second-order transport coefficients are related to the first-order ones through relations derived in kinetic theory. The one exception is the shear relaxation time τ_π , whose normalization is allowed to vary in a wide range. This allows for a precise quantification of the effect of a second-order transport coefficient on phenomenological constraints on η/s and ζ/s .

Figure 12 examines the extent to which the shear relaxation time normalization factor b_π affects the posterior of first-order transport coefficients. When trying to fit the experimental data, we see that in general smaller values of the shear relaxation time lead to larger shear viscosities, and vice versa. This is because increasing either b_π or η/s tends to reduce the harmonic flows, for instance.

Some sensitivity to the shear relaxation time factor b_π may be caused by the use of free-streaming as a pre-hydrodynamic model. Indeed, free-streaming can generate large values of $\pi^{\mu\nu}$. The subsequent size of $\pi^{\mu\nu}$ in hydrodynamics, following free-streaming, is governed by the overall size of τ_π which controls the relaxation towards the Navier-Stokes limit $\pi_{NS}^{\mu\nu} = 2\eta\sigma^{\mu\nu}$. Previous viscous hydrodynamics studies, using different initial conditions (not including free-streaming), found a smaller sensitivity to b_π [152, 153]. Note however that these studies did not include higher harmonic flows (v_3, v_4, \dots) which we find to have a stronger sensitivity to b_π than v_2 . This will be discussed in Section X.

Since there is significant uncertainty induced by the shear relaxation time on η/s (and to a lesser extent on ζ/s), future efforts should consider varying other second-order transport coefficients. Among those, the most important is likely the bulk relaxation time. Performing a systematic analysis of all second-order transport coefficients would be a significant future undertaking, in part because their parametric dependence must be specified, and a prior needs to be fixed before performing parameter estimation.²³ Non-conformal hy-

²³ Existing microscopic calculations such as Ref. [80] can help constrain the

parameter	full prior range	restricted prior range or value
α	$[-0.3, 0.3]$	0.0
T_η [GeV]	$[0.13, 0.3]$	0.154
a_{low} [GeV $^{-1}$]	$[-2, 1]$	0.0
λ_ζ	$[-0.8, 0.8]$	0
b_π	$[2, 8]$	5
p	$[-0.7, 0.7]$	$[-0.5, 0.5]$
w [fm]	$[0.5, 1.5]$	$[0.5, 1.0]$
τ_R [fm/c]	$[0.3, 2]$	$[0.3, 1.5]$
$(\zeta/s)_{\text{max}}$	$[0, 0.25]$	$[0.01, 0.1]$
T_ζ [GeV]	$[0.12, 0.3]$	$[0.15, 0.2]$
w_ζ [GeV]	$[0.025, 0.15]$	$[0.025, 0.1]$

TABLE III. Table of full (left) and restricted (right) parameter ranges. The restricted prior is similar to the prior employed in Ref. [27].

drodynamic has a large number of second-order transport coefficients [108]; relatively little is known for many of them. There is value in simultaneously studying these transport coefficients theoretically and phenomenologically. Even if these coefficients cannot be constrained from measurements, their influence on other model parameters should be studied.

IX. SENSITIVITY TO PRIOR KNOWLEDGE AND ASSUMPTIONS

Each model parameter in this study is assumed to have a uniform prior probability density over a finite range. This range represents crucial prior knowledge or assumptions. Our wider, less subjective prior (Table III, middle column) almost completely²⁴ encloses as a subspace the narrower, more subjective prior range postulated in Ref. [27] (Table III, right column).

By comparing the posteriors for these different prior parameter ranges we assess the sensitivity of our inference to prior knowledge. This is illustrated in Fig. 13 for one of the particlization models studied in this work (the Pratt-Torrieri-Bernhard model [53, 112]). We compare the posteriors for the specific shear and bulk viscosities using either the more or less subjective priors described above. These posteriors were obtained via Bayesian parameter estimation using only the ALICE Pb-Pb measurements at $\sqrt{s_{\text{NN}}} = 2.76$ TeV.

Clearly the more subjective prior drastically reduces the width of the credible intervals of the posterior for both the shear and bulk viscosities. Table III shows that, in addition to narrower prior ranges for the shear and bulk viscosities, the more restrictive prior assumed additional information about the initial conditions, and the shear relaxation time (which is a second-order transport coefficient).

parametric dependence and the priors. For example, Ref. [80] finds $\delta_{\pi\pi} = \frac{4}{3}\tau_\pi + \mathcal{O}((m/T)^2)$. One could assume $\delta_{\pi\pi} \propto \tau_\pi$ with a parameter being a proportionality factor of order 1.

²⁴ We do not include a curvature parameter for the shear viscosity at high temperatures as was done in Ref. [27] since there it was found, within the given prior limits, to be rather poorly constrained.

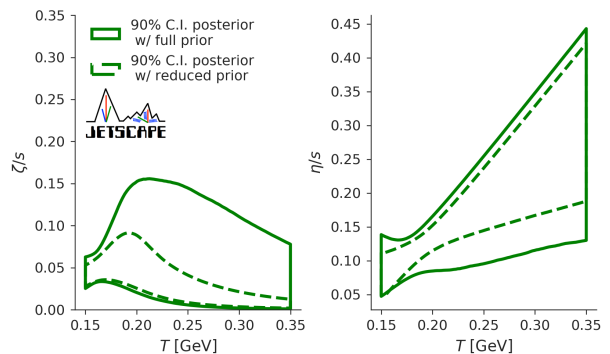


FIG. 13. The 90% posterior credible intervals of the specific bulk (left) and shear (right) viscosities for the Pratt-Torrieri-Bernhard viscous correction model, including only observables from LHC Pb-Pb collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV, depending on whether one uses a more informed or less informed prior.

Since the posterior of any Bayesian inference is proportional to the product of the prior and likelihood function, a tightening of the prior automatically also causes the posterior to tighten. Insofar the results shown in Fig. 13 are in principle expected. However, the observed large sensitivity of the posterior (in particular for the bulk viscosity) to the prior suggests that the constraining power of the experimental data is still limited, and that for a fully uninformed prior the 90% confidence intervals for the specific viscosities would be even wider than what is indicated by the solid lines in Fig. 13. Future improvements of the precision of our knowledge of the QGP viscosities require progress along at least one of the following two directions: (i) theoretical work leading to more objective priors, if not for the parameters of primary interest (i.e. the viscosities) then at least for the “nuisance parameters”; (ii) inclusion of additional measurements into the Bayesian analysis that have the potential to provide tighter likelihood functions for both the parameters of primary interest and the “nuisance parameters”. In the absence of theoretical progress towards tighter first-principles constraints on the viscosities, less subjective priors for these parameters of primary interest should be employed, to minimize sensitivity of the posterior to prior knowledge.

X. MODEL SENSITIVITY

To understand the posterior obtained by performing parameter estimation, it is useful to quantify which observables carry information about which model parameters. This can be achieved as follows:

- Select a set of parameter values at which the model sensitivity is explored
- Vary a single parameter at a time and quantify how much each observable responds

Note that the model is non-linear, and consequently this is a *local* measure of observables sensitivity, at a given point in the multidimensional parameter space.




FIG. 14. Sensitivity indices for LHC observables measured in the 0 – 5% (left) and 40 – 50% (right) centrality bin (except for the mean p_T event fluctuation $\delta p_T/p_T$ for which the 40 – 45% bin is plotted on the right), as a function of all model parameters. Plotted in blue is the Grad viscous correction model, in red the Chapman-Enskog model, and in green the Pratt-Torrieri-Bernhard model. The bars show the sensitivity to a 10% change in each parameter ($\delta = 0.1$).

Following Ref. [154] we define a local sensitivity index as follows: define two points in parameter space by $\mathbf{x} = (x_1, x_2, \dots, x_j, \dots, x_p)$ and $\mathbf{x}' = (x_1, x_2, \dots, (1 + \delta)x_j, \dots, x_p)$ where δ is a fixed percent difference. We use our emulator to predict all of the observables at these two points in parameter space. Suppose for some particular observable O , the emula-

tor predicts $\hat{O} = \hat{O}(\mathbf{x})$. Then, defining the percent difference in the observable by

$$\Delta \equiv \frac{\hat{O}(\mathbf{x}') - \hat{O}(\mathbf{x})}{\hat{O}(\mathbf{x})}, \quad (58)$$

our “sensitivity index” $S[x_j]$ for observable O under a change in parameter x_j is given by

$$S[x_j] \equiv \Delta/\delta. \quad (59)$$

We chose \mathbf{x} to be defined as the average of the three different Maximum A Posteriori (MAP) parameters (see Section V E) of each three viscous correction models, listed in Table II.

These sensitivity indices $S[x_j]$ for pairs of observables and parameters are shown in Fig. 14 for select Pb-Pb observables at $\sqrt{s_{\text{NN}}} = 2.76$ TeV and a step size $\delta = 0.1$. We verified that we obtain quantitatively similar results with a larger parameter step size $\delta = 0.4$, indicating that the parameter dependence of the model is reasonably close to linear in the region of parameter space studied. Note that the inclusion of emulator uncertainties in the sensitivity analysis is left as future work.

Although a local measure of the response of the model observables to changes in parameters is a strong approximation, it can nonetheless help guide our understanding regarding which observables carry information about each of the parameters. First note that the scale of the sensitivity indices is different for each parameter. Changing the shear relaxation time normalization b_π has a very small effect on all observables investigated in this work, with a 10% change in b_π leading to less than 1% change in observables.

On the other hand, very strong dependence on the model parameters have been found. The proton yield shows strong sensitivity to the switching temperature. Increasing the switching temperature by 10% increases the proton yield by about 20%. Consequently, most of the constraining power (or information) about the switching temperature is carried by the proton yield among the observables used herein.

As noted throughout this study, many of the observables show stronger sensitivity to the maximum of the specific bulk viscosity when the Pratt-Torrieri-Benhard distribution was employed, compared to any other viscous corrections used here. Looking again at Fig. 9, the strong sensitivity of the Pratt-Torrieri-Benhard viscous correction causes the 90% posterior credible interval for the specific bulk viscosity to be most tightly constrained among the viscous correction models explored here. The narrower posterior of ζ/s for this viscous correction model is a direct consequence of these larger sensitivity indices.

The parameter w in T_RENTo is largely responsible for controlling the eccentricities of the initial state. We find that the elliptic, triangular and quadrangular flows $v_2\{2\}$, $v_3\{2\}$, and $v_4\{2\}$ show strongest sensitivity among the observables plotted in Fig. 14. This may be expected from hydrodynamic response, in which $v_2\{2\} \propto \epsilon_2$, $v_3\{2\} \propto \epsilon_3$ and $v_4\{2\} \propto v_2\{2\}^2$. In addition, the initial geometry is more sensitive to the nucleon width w for peripheral collisions: we see that the harmonic flows for 40–50% centrality bins show close to twice the sensitivity to the width parameter than for 0–5% centrality.

As discussed in Section X, the triangular flow $v_3\{2\}$ and quadrangular flow $v_4\{2\}$ show strongest sensitivity to b_π in our model. This sensitivity remains small however: a 10% change in the shear relaxation time leads to a 1% change in $v_4\{2\}$ for example. This explains the challenge of constraining the shear relaxation time.

We highlight that in addition to guiding our understanding of the model, sensitivity measures can also direct where future experimental efforts can be focused [24] in order to best constrain model parameters, which is of interest to the entire community. We hope to expand our efforts in this direction in future studies with global (rather than local) sensitivity measures.

XI. BAYESIAN MODEL SELECTION

In addition to the estimation of parameters of a given model, Bayesian inference can also be used to quantify which model, among several competing models, is better supported by the experimental observations. There exist several metrics for comparison; we choose to employ the Bayes factor. We will first illustrate the application of the Bayes factor toward three of the viscous correction models for particlization that were used throughout this work. We then use it to compare our hydrodynamic model with simpler models which are ‘nested’ inside. Finally, the Bayes factor is applied towards answering whether a consistent model, with the same set of parameters describing the system created in RHIC Au-Au and LHC Pb-Pb collisions, or more complicated models where some parameters are allowed to differ, are better justified in light of the experimental data.

In the context of model selection, a ‘model’ refers to a specific set of parameters together with their prior, and a unique map from the parameters to a set of observables. As an example, a polynomial of second degree is a different model than a third-degree polynomial. In this case, however, the second degree polynomial model is ‘nested’ inside the polynomial of third degree; when the coefficient of the cubic term is fixed to zero, we can recover the second-order polynomial model. The polynomial of third degree has additional model complexity, given by the additional parameter and its prior. On the other hand, we can also compare a model given by a second degree polynomial with for example a model given by a sinusoid, with an uncertain amplitude and phase velocity. These models give altogether different predictions, and their coefficients have different meanings. In both cases above, whether the models share similar features or not, we will refer to them as different models when comparing them with the Bayes factor (which is the standard terminology). If one of the models happens to be nested inside the other, we will make note.

It may be useful to provide a word of caution at this point, as the use of the Bayes factor for comparing models has been met with some skepticism in the past; this skepticism is not necessarily undue. It is well known that the Bayes factor becomes ill-defined when one or both models have an ‘improper’ (non-normalizable) prior. In addition, the marginal evidence of each model in principle depends on the priors that were chosen. Neither of these potential issues is a roadblock for our current purposes, however. The priors for all model parameters considered are not improper, and have been purposely selected, weighing the applicability of each theoretical model along with reasonable information stemming from other sources of experimental data. A more thorough elucidation of

the Bayesian methodology, parameter estimation, model selection, and examples in cosmology can be found in [155].

A. Overview and estimation of Bayes factors

1. Bayes factor definition and interpretation

The Bayes factor is a useful measure for evaluating the relative merit of two competing models A and B in light of a given set of experimental data \mathbf{y}_{exp} . It is the ratio of the conditional probabilities, or the ‘odds’,

$$B_{A/B} \equiv \frac{\mathcal{P}(A|\mathbf{y}_{\text{exp}})}{\mathcal{P}(B|\mathbf{y}_{\text{exp}})} = \frac{\mathcal{P}(\mathbf{y}_{\text{exp}}|A) \mathcal{P}(A)}{\mathcal{P}(\mathbf{y}_{\text{exp}}|B) \mathcal{P}(B)}, \quad (60)$$

where Bayes’ theorem was applied to both the numerator and denominator. The odds depend on both the ratio of the likelihoods, marginalized over all parameters, multiplied by the ratio of prior beliefs about the validity of these two models. Unless there is a strong reason to believe that one model is much more likely than another, the ratio of priors is taken to be unity. In this case, the odds are simply the ratio of the Bayesian evidences of the two models:

$$B_{A/B} = \frac{\mathcal{P}(\mathbf{y}_{\text{exp}}|A)}{\mathcal{P}(\mathbf{y}_{\text{exp}}|B)}. \quad (61)$$

Using the marginalization and product rules for probabilities we ‘integrate in’ the model parameters \mathbf{x}_A for model A ,

$$\mathcal{P}(\mathbf{y}_{\text{exp}}|A) = \int d\mathbf{x}_A \mathcal{P}(\mathbf{y}_{\text{exp}}|\mathbf{x}_A, A) \mathcal{P}(\mathbf{x}_A) \quad (62)$$

and likewise for model B . The integrand appearing in Eq. (62) is the product of the same likelihood and prior which have been discussed in section V A. Therefore, the Bayesian evidence is simply the average of the likelihood with respect to the prior probability density. We note that when using priors which are uniform distributions on a finite domain, this expression can be simplified, yielding

$$\mathcal{P}(\mathbf{y}_{\text{exp}}|A) = \frac{1}{\mathcal{V}_A} \int_{\mathcal{D}_A} d\mathbf{x}_A \mathcal{P}(\mathbf{y}_{\text{exp}}|\mathbf{x}_A, A), \quad (63)$$

where we have defined the total volume of the prior \mathcal{V}_A for model A . This is the volume of the hypercube \mathcal{D}_A inside which the prior for model A is nonzero.

Because all of the models for which we perform model comparisons have uniform priors, the interpretations of our results are more straightforward. The Bayesian evidence of each model is the integral over the likelihood for the model inside the prior bounds, divided by the volume of the prior. Belief in a model is increased by ability to fit the data (larger likelihood), averaged inside of the prior bounds. But belief in the model is decreased by its ‘complexity’ (the ‘Occam penalty’), which is the volume of its prior which is excluded by the data. In a situation where the likelihood $\mathcal{P}(\mathbf{y}_{\text{exp}}|\mathbf{x}_A, A)$ does not actually depend on a particular model parameter $x_{A,i}$, we see from Eq. (63) that there is no Occam penalty because the size

of the prior region cancels in the numerator and denominator. Therefore, the Bayes factor does not penalize a model for having a parameter which is unconstrained by the data. For a more thorough explanation with examples, see Ref. [155].

2. Numerical methods for estimating the Bayes evidence

The integral over all the model parameters is very high-dimensional and does not lend itself to elementary methods. Fortunately, there exist methods for estimating the evidence that are ready to use in the existing Markov Chain Monte Carlo implementation [156] used throughout this work. A ‘parallel-tempered’ Markov Chain Monte Carlo routine defines a ladder of inverse ‘temperatures’ β_i , and then evolves an ensemble of walkers by sampling from distributions defined by

$$[\mathcal{P}(\mathbf{y}_{\text{exp}}|\mathbf{x}_A, A)]^{\beta_i} \mathcal{P}(\mathbf{x}_A). \quad (64)$$

We see that in the limit $\beta \rightarrow 0$, we recover our prior $\mathcal{P}(\mathbf{x}_A)$. At regular intervals each walker inside of each tempered distribution has the opportunity to swap positions with walkers at adjacent temperatures. Walkers at very high temperatures are not strongly affected by peaks in the likelihood function, while walkers at $\beta = 1$ are sampling from the target posterior. This gives this algorithm the advantage that it can efficiently sample multimodal distributions, which can be more difficult for other algorithms, including the ordinary Metropolis-Hastings, to sample accurately. Besides these advantages, the ladder of tempered distributions also gives an estimation of the Bayes evidence by the following trick. Defining the Bayesian evidence as a function of inverse temperature:

$$Z(\beta) = \int d\mathbf{x}_A [\mathcal{P}(\mathbf{y}_{\text{exp}}|\mathbf{x}_A, A)]^\beta \mathcal{P}(\mathbf{x}_A) \quad (65)$$

we note that it satisfies a differential equation

$$\begin{aligned} & \frac{d \ln Z}{d\beta} \\ &= \frac{1}{Z(\beta)} \int d\mathbf{x}_A \mathcal{P}(\mathbf{x}_A) \ln[\mathcal{P}(\mathbf{y}_{\text{exp}}|\mathbf{x}_A, A)] [\mathcal{P}(\mathbf{y}_{\text{exp}}|\mathbf{x}_A, A)]^\beta \\ &\equiv \langle \ln [\mathcal{P}(\mathbf{y}_{\text{exp}}|\mathbf{x}_A, A)] \rangle_\beta. \end{aligned} \quad (66)$$

Therefore, $\ln Z(\beta = 1)$ can be estimated by integrating by quadrature the average at each temperature. The uncertainty in this estimate $\delta \ln Z$ is primarily from using a finite number of points in the quadrature (finite number of ‘temperatures’).

B. Comparing viscous correction models

As a first illustration of Bayesian model selection, we quantify if our experimental data give evidence to prefer one viscous correction model over another. We have estimated the logarithm of the Bayes evidence $\ln Z$ as well as the integration uncertainty $\delta \ln Z$ for three of the four models using the parallel-tempering described above. Their mutual Bayes factors are shown in Table IV.

Model A	Model B	$\ln B_{A/B}$
Grad	CE	8.2 ± 2.3
Grad	PTB	1.4 ± 2.5
PTB	CE	6.8 ± 2.4

TABLE IV. A table of the logarithm of the Bayes factor $\ln B_{A/B}$ for each pair of viscous correction models and its integration uncertainty for the Grad, Chapman-Enskog (CE) and Pratt-Torrieri-Bernhard (PTB) viscous correction models.

From the Table we see that the Grad and Pratt-Torrieri-Bernhard models have Bayesian evidences that are compatible within the numerical uncertainty. The odds that the Grad model is better than the Pratt-Torrieri-Bernhard model are about 3:1, given our 0.6σ observation.²⁵ Therefore, the p_T -integrated calibration observables cannot distinguish which of these two models is more likely. However, we have moderate evidence to conclude that both of these models work better to describe the hadronic observables studied in this work than the Chapman-Enskog model, with the Grad versus Chapman-Enskog comparison being a 3.6σ observation (odds about 5000:1), and Pratt-Torrieri-Bernhard comparison a 2.8σ observation (odds about 400:1).

For the Chapman-Enskog model, we note from Fig. 10 that the marginal posterior of the free-streaming energy dependence α has a local maximum for $\alpha \lesssim -0.3$. It is possible that widening our prior to include smaller values of α would also increase the Bayes evidence for the Chapman-Enskog model. Unfortunately, this would require a new set of model calculations at new design points, which is beyond the scope of the present work. Due to the very large odds of the Grad model compared to the Chapman-Enskog model given by the Bayes factor, we also considered the frequentist odds defined by the maximum likelihood ratio. If L_A is the maximum value of the likelihood function for model A, and L_B the same for model B, this maximum likelihood ratio is simply defined by L_A/L_B . These maximum-likelihood odds were found to be close to 300:1 for the ratio of Grad to Chapman-Enskog models.

The Chapman-Enskog model is not able to simultaneously fit the proton multiplicity together with the other observables, such as the pion multiplicity. This puts the model under tension, and reduces the likelihood (averaged across the parameter space) of the Chapman-Enskog model. This is illustrated by Fig. 15, which displays the single and joint posterior predictive distributions of select observables for the most central bin 0–5% for Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV. For each of the Grad model (blue) and Chapman-Enskog model (red), parameter samples are drawn from the posteriors calibrated to all observables at both LHC and RHIC. Then, the model prediction is calculated using the emulator for all observables, and plotted is the model-experiment discrepancy, i.e., the difference between model prediction and experimental mean normalized by experimental standard deviation. That the chemi-

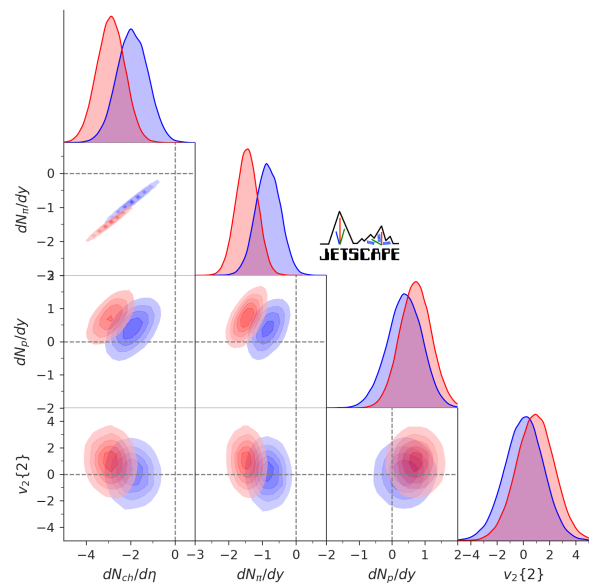


FIG. 15. Diagonal and off-diagonal panels show one- and two-dimensional projections of the n -dimensional posterior predictive distributions for selected Pb-Pb observables at fixed collision centrality of 0–5%. Plotted are the discrepancies between prediction and measurements in units of the experimental standard deviation; axes are labeled with shorthand notation $y \equiv (y_{\text{model}} - y_{\text{exp}}) / \sigma_{\text{exp}}$ where y stands for the observable whose model discrepancy is shown. The Grad model is shown in blue and Chapman-Enskog in red.

cal abundances disfavor the Chapman-Enskog model was further strengthened by recalculating the posteriors and Bayes factor for the Grad and Chapman-Enskog models excluding the LHC proton multiplicity from the data. In this case, the odds were greatly reduced to only about 5:1 in favor of Grad. For both the linearized Grad and Chapman-Enskog models, it is the bulk viscous correction which changes the chemical abundances from their equilibrium values (the shear viscous correction does not correct the equilibrium yields). Therefore, in light of the chemical abundances being a strong discriminator, it is specifically the bulk viscous correction given by the Chapman-Enskog model which is disfavored by the particle yields.

In conclusion, the hadronic observables studied in this work favor the Grad and Pratt-Torrieri-Bernhard models of viscous corrections over the Chapman-Enskog model. This is due in large part because the Chapman-Enskog model is worse at simultaneously fitting the chemical abundances. In light of the caveats, we do not believe this finding should be taken as a blanket statement on the validity of the Chapman-Enskog model in studying heavy ion collisions. Future studies will be necessary to clarify if viscous corrections can be systematically constrained from measurements.

Knowing the relative odds between the different particlization models, a model-averaged posterior with improved uncertainties for the inferred parameters can be derived using Bayesian Model Averaging which calculates the posterior as a weighted average of the individual model posteriors, weighted

²⁵ The probability is given by the left-tailed p -value.

by the Bayes evidence for each model. We reported such a result in Ref. [28].

C. Comparing hydrodynamic models

As another application of Bayesian model selection, we quantify whether simpler models, which are nested within the model described in Section III, are favored or disfavored by the data. We will make comparisons with models with simplified assumptions for the shear viscosity. As a reminder, the more complex model will be penalized by the additional parameters (the ‘Occam penalty’) that are constrained by the data, and will only yield a larger evidence than the simpler model if the extra constrained parameters significantly improve the fit to the data. An additional model parameter that is not well-constrained by the data within the range of the prior will have an insignificant Occam penalty.

1. Temperature independent specific shear viscosity

We consider whether our full model, which includes a temperature-dependent specific shear viscosity, is preferred by the data to a simpler model with a temperature-independent specific shear viscosity. In both cases we use the Grad viscous correction. We denote by model A our usual model with temperature dependent specific shear viscosity. We denote by B the model in which the low-temperature and high-temperature slopes a_{low} and a_{high} are fixed to zero; the temperature of the kink T_η is irrelevant in this scenario, and is also fixed to an arbitrary value. We find the logarithm of the Bayes factor to be consistent with zero within its uncertainty, $\ln B_{A/B} = -0.2 \pm 2.4$. Hence, the selected experimental data provide no evidence in favor of the common theoretical preference for a temperature-dependent specific shear viscosity of QCD matter. As noted above, the Occam penalty for including the additional parameters, which here are the slopes of the specific shear viscosity and position of its inflection, is minimal; this is because these parameters are not well constrained within the range of the prior. In any case, this “null” result suggests we should include more discriminating observables in future studies, beyond the subset we have chosen here.

2. Zero specific shear viscosity

We also study if our data provide strong evidence for a nonzero specific shear viscosity. This can be quantified in the same way as above, setting the parameters for the specific shear viscosity such that $(\eta/s)(T) \equiv 0$. We again use the Grad viscous correction model for this comparison and allow the specific bulk viscosity (as well as all other parameters) within their full prior ranges. We find the logarithm of the Bayes factor $\ln B_{A/B} = 11.7 \pm 2.6$ where model A is the default model with nonzero specific shear viscosity while model B has $\eta/s \equiv 0$. We conclude that the data provide strong evidence that the hy-

drodynamic models with nonzero specific shear viscosity are preferred.

D. Quantifying tension between LHC and RHIC

The Bayes factor is also useful for quantifying if our model is under significant tension when trying to simultaneously fit the data at both collision energies [157]. Throughout this work we have assumed that all model parameters are shared between the two systems except for their initial energy density normalizations. We can however relax these assumptions, and allow parameters to be different for the two different systems.

1. No common parameters between collision systems

Suppose that we allow all of the parameters to be different for the two systems defined by collisions at RHIC and LHC, respectively, including the initial conditions, viscosities, and switching temperature. In this case, our model has a total of 32 parameters. We want to compute the Bayes factor $\ln B_{A/B}$ where model A is the default model while model B assumes independent sets of model parameters for describing the data collected at different collision energies.

As usual, we take the ratio of our prior beliefs about these two models to be unity, $\mathcal{P}(A)/\mathcal{P}(B) = 1$, such that the Bayes factor reduces to the ratio of marginal evidences:

$$B_{A/B} = \frac{\mathcal{P}(\mathbf{y}_{\text{LHC}}, \mathbf{y}_{\text{RHIC}}|A)}{\mathcal{P}(\mathbf{y}_{\text{LHC}}, \mathbf{y}_{\text{RHIC}}|B)}. \quad (67)$$

As a consequence of the assumed statistical independence of measurements performed with different detectors at different collision energies, we can estimate the model evidence in the denominator as follows:

$$\mathcal{P}(\mathbf{y}_{\text{LHC}}, \mathbf{y}_{\text{RHIC}}|B) = \mathcal{P}(\mathbf{y}_{\text{LHC}}|B)\mathcal{P}(\mathbf{y}_{\text{RHIC}}|B). \quad (68)$$

Integrating over the model parameters for the LHC model yields

$$\begin{aligned} \mathcal{P}(\mathbf{y}_{\text{LHC}}|B) \\ = \int d\mathbf{x}_{\text{LHC}} \mathcal{P}(\mathbf{y}_{\text{LHC}}|B, \mathbf{x}_{\text{LHC}})\mathcal{P}(\mathbf{x}_{\text{LHC}}|B) \end{aligned} \quad (69)$$

which we have estimated using Eq. (66). A similar result holds for the model describing the RHIC data.

In this way we find $\ln B_{A/B} = 24.1 \pm 2.6$. We conclude that our data from the LHC and RHIC give very strong evidence that a model in which all parameters except the initial energy density normalizations are the same is strongly preferred by the data over a model in which all parameters are allowed to be different. The Occam penalty for basically doubling the number of model parameters far outweighs the gain of precision in the data description. We take this as strong evidence that a hybrid viscous hydrodynamic model with a single set of parameters provides a coherent physics picture for the experimental data measured at two collision energies that differ by over an order of magnitude.




FIG. 16. Partial representation of the posterior for a model that uses the Grad particlization model and allows for different nucleon width parameters w at RHIC and LHC energies. The estimated nucleon widths at the two collision energies inferred from the Bayesian analysis are found to agree within the 90% confidence limits.

Admittedly, allowing all of the parameters to be different leads to a very drastic comparison, adding far more model complexity than perhaps reasonable, thus entailing an outsized Occam penalty. A more meaningful study might have tried to identify tensions between a few specific observations and their predictions from the calibrated model, use the sensitivity analysis to zero in on one or a small number of model parameters to which these data are most sensitive, and introduce a small number of extra parameters to introduce additional collision energy dependence in that sector, with the aim of reducing the tensions. We leave this for future work.

2. Different transverse length scales in the initial conditions

We mentioned earlier that some theoretical models of the energy deposition in a heavy ion collision feature transverse length scales that depend on the collision energy. In our T_RENTO model, it is the “nucleon width” w which controls the transverse length scale for fluctuations in the initial conditions, and we have so far assumed that its value is independent of the collision energy. To test this assumption, we calculate the posterior for a model that introduces one additional parameter to allow the nucleon width w to change from one collision energy to the other. The posterior for this model is shown in Fig. 16.

We see that the most probable value for the nucleon width $w[0.2 \text{ TeV}]$ in Au-Au collisions at RHIC is about 20% larger than the width $w[2.76 \text{ TeV}]$ in Pb-Pb collisions at the LHC, though both agree within uncertainty as shown in Fig. 16. We note that the Color Glass Condensate model predicts roughly a factor of two difference between the color flux tube diameters

at RHIC and LHC energies [127]. The measured total inelastic nucleon-nucleon cross section also increases by about a factor two from RHIC to LHC, indicating a possible growth of w by a factor $\sim \sqrt{2}$. If A denotes the default model and B the model where the nucleon widths at the two collision energies are allowed to differ, we find $\ln B_{A/B} = 0.7 \pm 2.5$. Within the uncertainty of the estimate, we can thus not distinguish which model is preferred. The amount of tension that is caused by ignoring energy dependence of the nucleon width is not significant, and if there is indeed a small gain in the model evidence due to a slightly improved description of the data it is erased by Occam’s penalty for the increase in model complexity.

XII. PREDICTING p_T -DIFFERENTIAL OBSERVABLES

From the Bayesian point of view, a model is more useful if it is capable of making accurate predictions for observables that were not used for its calibration. This fits the physicist’s frame of mind in which belief in a model’s veracity is increased when the model makes an accurate prediction of some observable (and similarly, models that make inaccurate predictions are held in lower esteem).

We should thus check whether our calibrated model for heavy-ion collision evolution makes accurate predictions. We consider as a prediction any observable calculated from the model using the Maximum A Posteriori (MAP) parameters (see Section V E and Table II) that has not been used for the model calibration, neither through a prior nor via the likelihood. As our model is intended to describe the physics of particles with soft momenta $p_T \lesssim 2 \text{ GeV}$, accurately predicted soft observables should increase our belief in the model, while soft observables that are inaccurately predicted will decrease it. As an example, in this section we use our model to predict the shapes of the p_T -differential identified hadron spectra and charged hadron elliptic flow measured by ALICE at the LHC, shown in Figs. 17 and 18 for the Grad and Chapman-Enskog particlization models.²⁶

Because the multiplicities and mean transverse momenta are dominated by particles with typical (flow-boosted) thermal momenta, the model tends to fit the slope of the pion differential spectra better at soft momenta $p_T \lesssim 1.5 \text{ GeV}$. The stronger boost from radial flow experienced by heavier hadrons [158, 159] extends this agreement with the model to higher $p_T \lesssim 2.5 \text{ GeV}$ for protons. This finding is consistent with that of Ref. [23] which showed that the shape of the pion and proton spectra could be characterized very well by the mean transverse momenta and yields.

For the differential elliptic flow, the agreement between model prediction and experiment is generally good for both the Grad and Chapman-Enskog models; neither model performs qualitatively better than the other. To what extent each of these

²⁶ We remind the reader that the posterior of our model parameters was estimated using only p_T -integrated observables, e.g. the multiplicities and mean transverse momenta for pions, kaons and protons, the p_T -integrated harmonic flows, etc.




FIG. 17. The transverse momentum spectra for pions (left), kaons (center) and protons (right) averaged over five thousand fluctuating events predicted by the Grad (dashed lines) and Chapman-Enskog (dotted lines) models, each run at their respective MAP parameters. Shown are the predictions for the 0 – 5% (red), 10 – 20% (green), 20 – 30% (blue), and 30 – 40% (cyan) centralities, each having been scaled by a power of five for visualization. Also shown are the measurements from ALICE (open circles). The bottom panel shows the ratio of the model prediction ‘JS’ divided by the ALICE data.

models’ predictions also agree with additional experimental results that were not used for model calibration will be further explored in future studies. We note that the Chapman-Enskog viscous correction model is not able to fit the experimental multiplicities of pions and protons as well as the Grad model, but in the p_T -differential elliptic flow the normalizations of these spectra drop out and only their shapes as a function of p_T matter.

XIII. SUMMARY

Building upon previous studies [22–27], this work presents a comprehensive framework to perform Bayesian inference in heavy-ion collisions. A major new addition compared to these earlier studies is the use of closure tests (Section VI), which we used to perform an extensive validation of our analysis before comparison to data. We further discussed how closure tests can be used to explore the constraining power of different observables, before measurements are even performed, thus helping prioritize which observables need better measurements.

We highlighted the important role played by the parameter “priors” in Section IV. The prior probability distribution encodes how likely we believe the model parameters to take certain values, before comparison with data. Failure to quantify the impact of the priors on the posterior can lead to incorrect or misleading conclusions about the constraining power of heavy-ion measurements.

Using the above Bayesian framework, duly validated with closure tests, a parameter estimation was performed using RHIC and LHC data, first using measurements from these two colliders separately and then combining them (Section VII). Our analysis shows that these RHIC and LHC data can set a strong constraint on the viscosity of hot nuclear matter in a temperature window between 150 and 200 MeV, which is

probed by a large fraction of the space-time volume filled by the expanding medium created in the collision. However, as hinted by closure tests (Section VI), we observed that the bulk viscosity of QCD is difficult to constrain for temperatures above ~ 200 MeV, at least with the hadronic observables used in this work. The shear viscosity was also found to be difficult to constrain for $T \gtrsim 250$ MeV.

A good simultaneous agreement with measurements from RHIC and the LHC was found, both when comparing with random samples from the parameter posterior (Fig. 7) and when comparing with the Maximum A Posteriori parameters (Fig. 8). Results from simulations utilizing the Maximum A Posteriori parameters furthermore agreed well with the p_T -differential spectra of identified hadrons (Fig. 17), and the p_T -differential v_2 of charged hadrons (Fig. 18).

All results presented in this work were performed with a new comprehensive simulation framework for heavy-ion collisions, described in Section III, which combines T_RENTo initial conditions, free-streaming, second-order relativistic hydrodynamics, a flexible particlization module and the recently developed SMASH hadronic transport as afterburner. The diversity of physics ingredients entails a large set of model parameters. The significance of these parameters was investigated by studying the sensitivity of predicted values for the experimental data to changes in these parameters (Section X). The Maximum A Posteriori (MAP) values for these parameters are listed in Table II, for three different particlization models that employ different parametrizations of viscous corrections to the momentum distributions and provide the mapping between the energy-momentum tensor of hydrodynamics and the momentum distributions of hadrons. We emphasize that the posterior probability distributions of the parameters, such as those in Fig. 9 and Fig. 10, contain much more information than just the Maximum A Posteriori. Nevertheless, for practical applications, a single set of model parameters must often be used.




FIG. 18. The p_T -differential two-particle cumulant elliptic flow of charged particles averaged over five thousand fluctuating events predicted by the Grad (blue) and Chapman-Enskog (red) models, run at their respective MAP parameters. Also shown are measurements from ALICE taking a pseudorapidity gap $\Delta\eta = 0.2$ (open circles) or $\Delta\eta = 1.0$ (filled circles)

For such applications, we put forward the Maximum A Posteriori parameters from Table II as a sensible choice.

The Maximum A Posteriori parameters from Table II take somewhat different values for the three different particlization models investigated herein. This is one of several examples of model uncertainty that were investigated in this work. Overall, these studies lend a considerable amount of credence to our constraints on the model parameters, such as the shear and bulk viscosities, and we expect them to provide valuable guidance for analyses that use somewhat different models of heavy-ion collisions.

This work makes clear that constraints on the viscosities of QCD still have substantial theoretical uncertainties from viscous corrections to the hadronic distributions. We showed in Fig. 9 that constraints on ζ/s and η/s can shift significantly when different viscous corrections are used. Not all model parameters were found to be equally sensitive to these viscous corrections. For example, the initial condition parameters from the T_RENTO ansatz do not generally depend heavily on the viscous corrections (Fig. 10). The free-streaming time,

on the other hand, was found to have a significant sensitivity (Fig. 11). This highlights the challenge of phenomenological constraints on properties of QCD: the details of how energy and momentum are distributed across species at the late particlization stage can have a significant effect on other model parameters describing much earlier evolution stages. The dependence of the free-streaming time on viscous corrections is significant. There have been associations made in the past between the free-streaming time τ_{fs} and the “hydrodynamization time” in heavy-ion collisions. We would like to express some wariness about this interpretation of τ_{fs} : our difficulty in constraining the centrality and center-of-mass energy dependence of the free-streaming time (Eq. (19) and Fig. 11) suggests that there are still significant model biases in our treatment of the initial stage of the collision.

The effect of the shear relaxation time on the inferred values for shear and bulk viscosity was quantified in Fig. 12. We believe this inclusion of a second-order transport coefficient in the model calibration to be a significant advance, given that significant theoretical uncertainties remain in the treatment of second-order transport coefficients in general (see the end of Section III B).

The result of our Bayesian parameter estimation is characterized by a non-trivial combination of the sensitivities of each observable to the model parameters. In Fig. 14, we showed explicitly the manner in which the observables react to changes in the values of the model parameters used in this work. As this dependence is non-linear, the local model sensitivity results are not universal. Nevertheless Fig. 14 provides valuable intuition on the contribution of different observables to constraining model parameters, connecting with previous works on the topic [24].

A further important tool explored here is Bayesian model selection (Section XI). Bayes factors were used to compare the level of tension with measurements of different viscous correction models (Section XI B). We found that chemical abundance measurements disfavor the Chapman-Enskog particlization model in comparison with the two other viscous correction models studied in this work; this begs for additional studies to more firmly assess the robustness of this conclusion.

Bayes factors reward improved model agreement with measurements and penalize increased model complexity. We used this feature in Section XI D 1 to verify that a joint description of both RHIC and LHC data by a single dynamical model with a common set of model parameters is favored over individual descriptions with separate sets of model parameter values. In Section XI D 2 we used the same approach to quantify the odds in favor of, or against, a collision energy dependence of the nucleon width parameter w in T_RENTO, finding no statistically significant evidence against using the same value at RHIC and LHC.

The entire model presented in this work, including the Bayesian inference tools employed for its calibration and the uncertainty quantification for the inferred model parameters will soon become publicly available as part of a new release of the JETSCAPE framework [32]. As described in Appendix H, we performed a careful and thorough validation of these new numerical implementations. An important outcome

of this validation exercise was a systematic comparison of the SMASH and UrQMD afterburners. We found that (at least for the model parameters used here) both afterburners produce very similar results for the hadronic observables studied in this work (Appendix H2). Importantly, we found that exact consistency of the particle species included in the equation of state and the hadron lists for the particlization and afterburner modules is essential for this test to succeed: an inconsistency in any of these ingredients could easily be misinterpreted as a difference in the afterburners themselves.

We believe that there is great value in the ability to easily (i.e. at little numerical cost) visualize how observables react to changes in individual or combinations of model parameters — a piece of information provided by the emulator. We have constructed a user-friendly “widget” that offers this feature for some of the LHC data used in this work and make it publicly available online in Ref. [160]. We believe this visualization tool can be of use to the heavy-ion community for better understanding the relation between the QGP viscosities and the hadronic observables measured experimentally, for example.

The present analysis builds on pioneering work published in several previous studies, in particular in Jonah Bernhard’s recent PhD thesis [27]. The results presented here differ from this latest analysis in several respects and for a number of reasons that are discussed in detail in Appendix G. We consider identifying the role of the parameter prior as particularly important. The smaller range of values for ζ/s explored in Ref. [27] could give the impression that good constraints were possible on the bulk viscosity at high temperature; however, by allowing ζ/s to take a wider range of values in this work, we observed that ζ/s is poorly constrained at temperature above 200 MeV.

The work reported here presents the state of the art in heavy-ion collision modeling using model calibration via Bayesian inference. The posterior probability distributions obtained here provide the most realistic and robust constraints available on key properties of the quark-gluon plasma created in relativistic heavy-ion collisions at RHIC and LHC. While the uncertainties quoted by us are quantified to the extent possible with the tools in our hands, they are still large for several parameters of primary interest. In the following Section we offer an outlook on future work that can perhaps help to further improve this situation.

XIV. OUTLOOK

There are many ways to build upon the results presented here. Within the model of heavy-ion collisions used in this work, one can proceed to include additional observables in the Bayesian analysis, to improve the still limited constraints on shear and bulk viscosity. Observables that could be especially interesting to include are the HBT radii, which are expected to provide complementary information to the development of radial flow and switching temperature [24, 161]. The inclusion of jet observables and electromagnetic probes in the Bayesian analysis, while a major undertaking, would represent milestones with considerable potential to constrain the viscosi-

ties at higher temperature as well as the properties of the early stage in heavy-ion collision [162–165]. The realization of the present study within the JETSCAPE Collaboration, using the JETSCAPE Framework, is meant as a step in this direction.

Evidently, the Bayesian analysis framework presented in this work can be extended to multiple collision systems with different sizes, from the ones at intermediate size such as LHC’s Xe-Xe collisions and RHIC’s isobar run (Ru-Ru and Zr-Zr), to smaller and asymmetric collision systems, such as RHIC’s p-Au, d-Au and He-Au collisions, LHC’s p-Pb and high-multiplicity p-p collisions, and O+O collisions at both colliders. Performing a Bayesian analysis with small systems has its challenges but also provides highly valuable insights of a systematic model-to-data comparison of collision systems of varying sizes using a single model [30, 31, 67].

As discussed throughout the manuscript, it is inevitable that the results of a Bayesian analysis are tied to the exact details of the model. For example, it is possible that the inferred values of the shear and bulk viscosities of QCD would change non-trivially if different models of pre-hydrodynamic physics were used. The flexibility of our $T_{\text{RENT}} + \text{free-streaming}$ ansatz captures much of the uncertainty from the pre-hydrodynamic phase, and folds this uncertainty into our posteriors for the viscosities of QCD. Nevertheless we do believe that the pre-hydrodynamic stage presents one of the major remaining sources of uncertainty in constraining the viscosities of QCD from heavy ion measurements; it will be important to explore different pre-hydrodynamic models in the future.²⁷

Another important source of uncertainty is expected to be the bulk relaxation time. The current initialization of bulk pressure has known issues, and the bulk relaxation time will have a direct impact on the propagation of this initialization uncertainty onto the final hadronic observables. While we did not find a strong dependence on the bulk relaxation time for our hadronic observables, at least for our Maximum A Posteriori parameters (Appendix F), we do believe it is important for this second-order transport coefficients to be explored systematically in the future.²⁸

Finally, we note that this work includes no theoretical uncertainty from the equation of state. Such uncertainties may be larger than commonly expected [76] and should be explored in the future, building on works such as Refs [24, 76, 166].

These and other theoretical uncertainties, such as the mapping between hydrodynamics and the hadronic momentum distribution discussed in this work, currently limit our ability to constrain the QGP viscosities from heavy-ion collision measurements. Statistical methods for quantifying, within the framework of Bayesian inference, these modeling uncertainties are currently being developed; targeted strategies for further reducing them require additional theoretical effort.

²⁷ An example of Bayesian study with a different pre-hydrodynamic model that allows for breaking the assumption of conformal symmetry can be found in Refs. [30, 31]. In that model the expansion velocity is parametrized and can be used to control the sign and magnitude of the bulk pressure. This may improve the description of the pre-hydrodynamic phase.

²⁸ We note that Refs.[30, 31] include a first study of the bulk relaxation time using Bayesian inference.

ACKNOWLEDGMENTS

We thank Jonah Bernhard, Gabriel Denicol, Scott Moreland, Scott Pratt and Derek Teaney for useful discussions, and Richard J. Furnstahl and Xilin Zhang for their insights on Bayesian Inference and Markov Chain Monte Carlo. This work was supported in part by the National Science Foundation (NSF) within the framework of the JETSCAPE collaboration, under grant numbers ACI-1550172 (Y.C. and G.R.), ACI-1550221 (R.J.F., F.G., M.K. and B.K.), ACI-1550223 (D.E., M.M., U.H., L.D., and D.L.), ACI-1550225 (S.A.B., J.C., T.D., W.F., W.K., R.W., S.M., and Y.X.), ACI-1550228 (J.M., B.J., P.J., W.K., X.-N.W.), and ACI-1550300 (S.C., L.C., A.K., A.M., C.N., A.S., J.P., L.S., C.Si., R.A.S. and G.V.). It was also supported in part by the NSF under grant numbers PHY-1516590, 1812431 and PHY-2012922 (R.J.F., B.K., F.G., M.K., and C.S.), and by the US Department of Energy, Office of Science, Office of Nuclear Physics under grant numbers DE-AC02-05CH11231 (D.O., X.-N.W.), DE-AC52-07NA27344 (A.A., R.A.S.), DE-SC0013460 (S.C., A.K., A.M., C.S. and C.Si.), DE-SC0004286 (L.D., M.M., D.E., U.H. and D.L.), DE-SC0012704 (B.S. and C.S.), DE-FG02-92ER40713 (J.P.) and DE-FG02-05ER41367 (T.D., W.K., J.-F.P., S.A.B. and Y.X.). The work was also supported in part by the National Science Foundation of China (NSFC) under grant numbers 11935007, 11861131009 and 11890714 (Y.H. and X.-N.W.), by the Natural Sciences and Engineering Research Council of Canada (C.G., M.H., S.J., C.P. and G.V.), by the Fonds de recherche du Québec – Nature et technologies (FRQNT) (G.V.), by the Office of the Vice President for Research (OVPR) at Wayne State University (Y.T.), by the São Paulo Research Foundation (FAPESP) under projects 2016/24029-6, 2017/05685-2 and 2018/24720-6 (M.L.), and by the University of California, Berkeley - Central China Normal University Collaboration Grant (W.K.). U.H. would like to acknowledge support by the Alexander von Humboldt Foundation through a Humboldt Research Award. Allocation of supercomputing resources (Project: PHY180035) were obtained in part through the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number ACI-1548562. Calculation were performed in part on Stampede2 compute nodes, generously funded by the National Science Foundation (NSF) through award ACI-1134872, within the Texas Advanced Computing Center (TACC) at the University of Texas at Austin [167], and in part on the Ohio Supercomputer [168] (Project PAS0254). Computations were also carried out on the Wayne State Grid funded by the Wayne State OVPR, and on the supercomputer *Guillimin* from McGill University, managed by Calcul Québec and Compute Canada. The operation of the supercomputer *Guillimin* is funded by the Canada Foundation for Innovation (CFI), NanoQuébec, Réseau de Médecine Génétique Appliquée (RMGA) and FRQ-NT. Data storage was provided in part by the OSIRIS project supported by the National Science Foundation under grant number OAC-1541335.

Appendix A: Full posterior of model parameters

For completeness, we show in Fig. 19 the posterior of all model parameters single and joint-parameter marginal distributions for the Grad (blue) and Chapman-Enskog (red) viscous correction models, combining both RHIC and LHC experimental results.

Appendix B: Posterior for LHC and RHIC independently

In Fig. 20 we show the parameter estimates for select T_{RENT0} initial condition parameters, as well as the switching temperature. Each posterior was estimated using only observables from a single system; LHC observables (purple) or RHIC observables (orange).

In this study, we have included more observables for Pb-Pb $\sqrt{s_{\text{NN}}} = 2.76$ TeV collisions at LHC, so the likelihood functions are more tightly constrained, while those for Au-Au $\sqrt{s_{\text{NN}}} = 0.2$ TeV collisions at RHIC are broader. In addition, the switching temperature for RHIC posterior is poorly constrained because we have omitted the proton yield, which should be among the most sensitive observables. Overall, we see that there is good agreement in the estimates of these parameters whether one uses Pb-Pb $\sqrt{s_{\text{NN}}} = 2.76$ TeV observables or Au-Au $\sqrt{s_{\text{NN}}} = 0.2$ TeV observables.

Appendix C: Validation of principal component analysis

Principal component analysis acts to identify the linear correlations among pairs of observables. A figure showing the correlations among all possible pairs of observables would be far too large to plot, but we plot a subset of possible pairs in Fig. 21 and make some important observations.

We see that certain pairs of observables have a strong linear correlation: for instance the yield dN/dy of pions in the 20–30% centrality bin and yield of charged particles $dN_{\text{ch}}/d\eta$ in the 30–40% centrality bin. For such pairs, nearly all the information about the model parameters is contained in just one of the observables. Uncorrelated pairs contain independent information about the model parameters. No pair of observables displays a significant non-linear correlation except for the elliptic flow $v_2\{2\}$ and quadrangular flow $v_4\{2\}$, which shows a correlation $v_4\{2\} \sim (v_2\{2\})^2$. The scarcity of strong non-linear correlations suggests that ordinary principal component analysis is a suitable method for dimensionality reduction. To test that our model emulator used for Bayesian inference is not overfit to features from statistical noise in the hybrid model, we have examined the effect on the posterior when we reduce the number of principal components by a factor of two for each system. In this case, five principal components explain about 94% of the variance of Pb-Pb collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV and three principal components about 91% of the variance of Au-Au collisions at $\sqrt{s_{\text{NN}}} = 200$ GeV data. The posteriors of the specific viscosities in these two cases are compared in Fig. 22.




FIG. 19. The posterior for Grad (blue) and Chapman-Enskog (red) viscous correction models for all model parameters, combining both RHIC and LHC experimental results. Units for dimensional quantities are those given in Table I.

The uncertainty contributed by the principal components that we omit contributes to the total emulator uncertainty and the posterior of specific shear and bulk viscosities is broadened in the case with fewer principal components included. To be sure that an emulator (and the choice of the number of principal components) is not underfit or overfit, one must perform emulator validation (see Section VI). An emulator which is overfit will fit the training points very well, but will perform

poorly in predicting the observables for a novel testing point.

Appendix D: Experimental covariance matrix

Currently, only the diagonal terms in the experimental covariance matrix are reported by the ALICE and STAR experiments. We have assumed a diagonal covariance matrix when

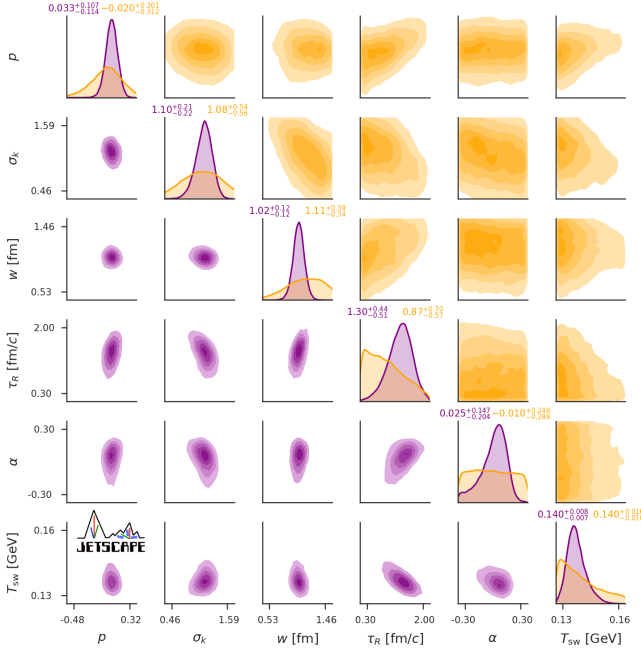


FIG. 20. The posterior of initial conditions and switching temperature for the Grad viscous correction model using only LHC data (purple) or only RHIC data (orange).

performing parameter estimation. However there are undeniably nontrivial correlations in the systematic uncertainties of measured observables and centrality bins. This is important, since systematic uncertainties are generally the dominant source, larger than statistical uncertainties. We test qualitatively how these correlated uncertainties may affect our analysis. The assumed covariance matrix will affect the posterior for all model parameters, but for simplicity we quantify its effect on the posteriors of specific shear and bulk viscosities. This is shown in Fig. 23 for the Grad viscous correction model.

In the case of the correlated experimental covariance matrix, given centrality bins c_i and c_j of the same observable, the experimental covariance is assumed to be

$$\Sigma_{i,j}^{\text{exp}} = \rho \sigma_i \sigma_j \quad (\text{D1})$$

where

$$\rho = \exp\left(-\frac{(c_i - c_j)^2}{l^2}\right) \quad (\text{D2})$$

and σ_i is the standard deviation of the observable in centrality bin c_i . Observables are organized in groups: (i) multiplicities; (ii) mean transverse momenta; (iii) harmonic flows; and (iv) transverse momentum fluctuations. For pairs of different observables within the same ‘group’, we take the same correlation coefficient defined above and multiply by an overall factor of 0.8. For pairs of different observables in different groups, we assume zero correlation. The ‘‘correlation length’’ between centrality bins is assumed $l = 0.5$ [27].

We see that an ansatz for the covariance matrix that includes nonzero correlations has a significant effect of broadening the viscous posterior, increasing the overall uncertainty. In the absence of a reported experimental covariance matrix,

a more Bayesian approach would be to treat the correlation length l and magnitude ρ as uncertain nuisance parameters in the Bayesian parameter estimation, with priors guided by the knowledge and study of the experimental collaborations, and marginalize over them. This is an important extension that we leave for future studies.

Appendix E: Reducing experimental uncertainty

We quantify the extent to which the experimental uncertainty contributes to the total uncertainty in our posterior for the specific shear and bulk viscosities. Besides experimental uncertainty, there is always nonzero uncertainty contributed by the use of a model emulator. We quantify this by changing artificially the uncertainty on experimentally measured observables during the parameter estimation; the result is shown in Fig. 24.

We see that significantly reducing the experimental error has the potential to qualitatively move our posterior for the specific bulk and shear viscosities. Perhaps more importantly, even if we reduce all of the experimental uncertainty by a factor of two, the credible intervals for the specific bulk and shear viscosities still remain quite large at high temperatures. This hints that in the future we should include additional observables and systems which are more sensitive to the viscosities at high temperatures.

Appendix F: Bulk relaxation time

Throughout this study, we have used a parametrization of the specific bulk viscosity given by

$$\tau_{\Pi} = b_{\Pi} \frac{\zeta}{\left(\frac{1}{3} - c_s^2\right)^2 (\epsilon + p)} \quad (\text{F1})$$

where $b_{\Pi} = 1/14.55$ [80]. We study how a change in b_{Π} translates into a change in our observables using the Maximum A Posteriori parameters for the Grad viscous correction model. This is shown in Fig. 25.

Because our pre-hydrodynamic model is free-streaming and conformal, the resulting bulk pressure at Landau matching is large and positive. The bulk pressure will relax to its Navier-Stokes value $\Pi_{\text{NS}} = -\zeta\theta$ on a time scale given by τ_{Π} . Increasing τ_{Π} we find that the bulk pressure stays positive for a longer time. Therefore, comparing the two sets of events, the calculations with the larger bulk relaxation time have larger mean transverse momenta and transverse energy. In future studies, it will be important to study to what extent the bulk relaxation time affects the posterior for the specific bulk and shear viscosities.

As the bulk relaxation time is further increased, our model also has the feature that the bulk pressure may not have time to relax to its Navier-Stokes value $\Pi_{\text{NS}} = -\zeta\theta$ during the lifetime of the hydrodynamic phase. In this case, the evolution of the bulk pressure becomes less sensitive to the value of the

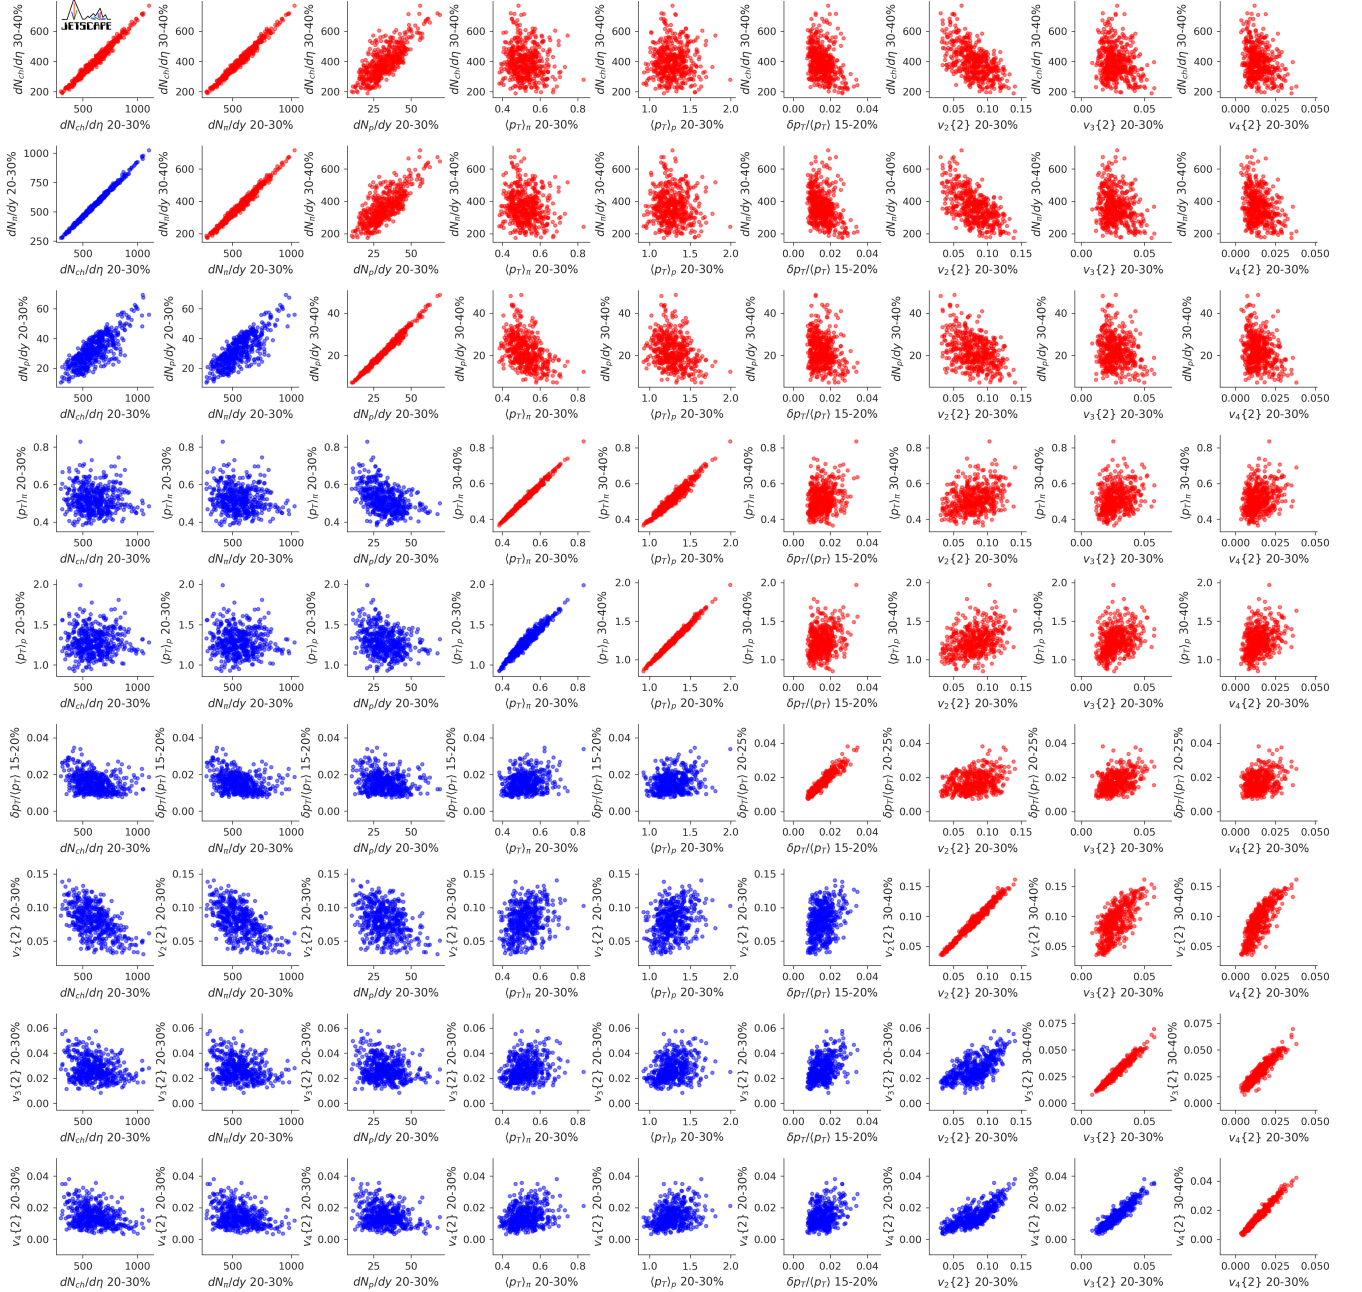


FIG. 21. Scatter plots of selected pairs of observables predicted by our model for Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV for the 500-point parameter design. Some pairs in the same centrality bin are shown in blue, while all pairs of different centrality bins are shown in red. Many pairs of observables have strong linear correlations, in which case they do not contain significant mutual information (knowing one is nearly sufficient). Pairs of observables which do not have strong linear correlations carry independent information about the parameters.

specific bulk viscosity and is dominated by its initial conditions. Inferring the likely values of the specific bulk viscosity then becomes more challenging.

Appendix G: Comparison to previous studies

In this section we enumerate the largest differences between the parameter estimation presented in this analysis and the analysis found in Ref. [27].

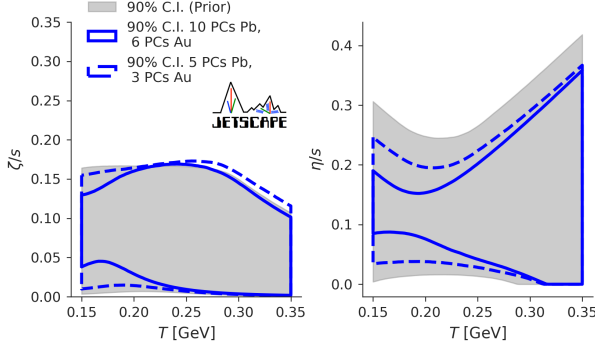


FIG. 22. Comparing the viscosity posteriors when we perform Bayesian parameter estimation with less principal components. The solid blue results from estimation using 10 principal components for the Pb-Pb $\sqrt{s_{NN}} = 2.76$ TeV emulator and 6 principal components for the Au-Au $\sqrt{s_{NN}} = 0.2$ TeV emulator. The dotted blue results from 5 principal components for Pb-Pb $\sqrt{s_{NN}} = 2.76$ TeV and 3 principal components for Au-Au $\sqrt{s_{NN}} = 0.2$ TeV.

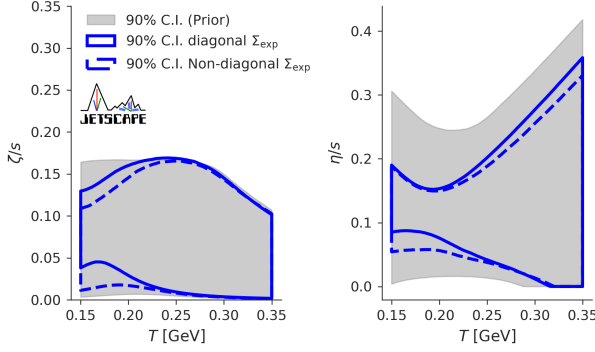


FIG. 23. The change in the viscous posterior resulting from assuming a diagonal experimental covariance matrix (solid blue band) or correlated experimental covariance matrix (dashed blue band). Details regarding the magnitude of correlations in text body.

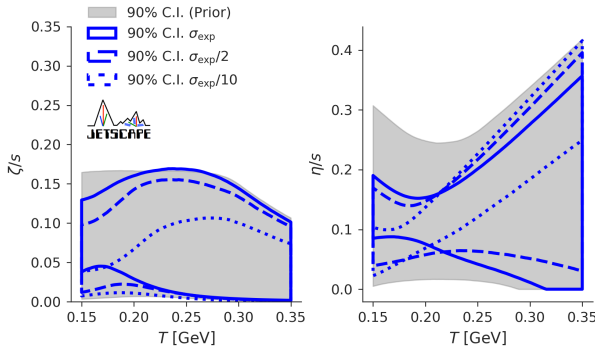


FIG. 24. The posterior for specific bulk (left) and shear (right) viscosities depending on whether it includes the full experimental uncertainties (filled blue) or divides them by a factor of 2 (dashed blue) or 10 (dotted blue). The model emulator always contributes non-zero uncertainty.




FIG. 25. The solid lines are averages over five thousand Pb-Pb $\sqrt{s_{NN}} = 2.76$ TeV events generated with the MAP parameters for the Grad model, and the default bulk relaxation time factor $b_{\Pi} = 1/14.55$. The circles are generated with the same set of parameters except $b_{\Pi} = 2/14.55$.

1. Physics models

a. Pre-hydrodynamic free-streaming: Both Ref. [27] and our own used free-streaming as a pre-hydrodynamic expansion model. Different numerical implementations were used, but they were validated against each other and found to be in excellent numerical agreement. However, in this work we

have allowed the free-streaming time to be dependent on the energy of each collision. This additional feature is manifest in the parameter α ; when α is fixed to zero, both studies have the same physics for the pre-hydrodynamic free-streaming.

b. Hydrodynamics: equation of state and viscosities: The largest differences in the hydrodynamic models include the equation of state and the parametrization of specific shear and bulk viscosities.

The equation of state used in Ref. [27] was given by the HotQCD lattice result at high temperatures matched to the 2017 PDG table of hadronic resonances at low temperatures. In particular, this included a very light σ meson with a mass of about 500 MeV. Our study has matched the same HotQCD lattice equation of state at high temperatures to a table of hadronic resonances entering in the SMASH afterburner. In particular, we excluded the σ meson entirely in the construction of the equation of state.

Besides the list of resonances which compose the hadron resonance gas component, Ref. [27] also computed the hadronic equation of state assuming relativistic Breit-Wigner resonances with nonzero width, while this study assumed all resonances on mass-shell in constructing the equation of state.

In the parametrization of the specific shear viscosity, Ref. [27] included a curvature parameter for the specific shear viscosity at high temperatures, which was not included in this work. On the other hand, we varied the slope of the low-temperature specific shear viscosity, as well as the position of the “kink” in this work, while both of these were fixed in Bernhard’s study. For the specific bulk viscosity, we allowed the parametrization to have a nonzero skewness, which was not present in Bernhard’s study.

c. Particlization, resonance width and σ resonance: Ref. [27] fixed the particlization model to be what we have referred to as the Pratt-Torrieri-Bernhard viscous correction model (Section III C), while in this work we also investigated other models.

For all viscous correction models in this work, the particles were sampled on their mass-shell, while particlization in Ref. [27] sampled the particles mass from a relativistic Breit-Wigner function.

In addition, as already mentioned, Ref. [27] sampled unstable σ resonances with a mass of about 500 MeV, which significantly increased the number of pions at low momenta once they decayed. This study excluded the σ resonance from sampling during particlization.

d. Hadronic afterburner: Finally, the hadronic afterburner used in Ref. [27] was UrQMD, while we use SMASH. Although these two models include somewhat different lists of resonances as well as slightly different hadronic cross-sections, we checked in Appendix H 2 that UrQMD and SMASH have excellent agreement when used with the model parameters that agree well with data. For that reason, we believe at this time that the difference in hadronic afterburners is negligible in comparison to the other differences listed above.

2. Prior distributions

The prior used in Ref. [27] is nearly a subspace of the prior used in this study, with the exception of the high-temperature behavior of the specific shear viscosity. Ref. [27] allowed the specific shear viscosity to have a nonzero curvature, i.e. quadratic temperature dependence at high temperatures. In this study, we have not allowed such a quadratic temperature dependence in the specific shear viscosity at high temperature.

3. Experimental data

Both Ref. [27] and this study have included the ALICE p_T -integrated, centrality-dependent data for Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV. However, Ref. [27] additionally included data for Pb-Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV, which are not included in this work. Instead, we have included STAR data for Au-Au collisions at $\sqrt{s_{NN}} = 200$ GeV, which were not included in Ref. [27].

Appendix H: Multistage model validation

Several of the numerical implementations of models used in this work are used for the first time; other required modifications and expansions. For this reason we include validations of these codes against counterparts which have been used extensively in previous studies.

1. Validation of second-order viscous hydrodynamics implementation

In this section, we compare two different numerical implementations of the same underlying second-order relativistic hydrodynamics equations [73]. The first implementation is the one used throughout this work is MUSIC [49–51]. The second implementation is a slightly modified version of the VISHNew 2+1D hydrodynamics code [74, 153], `osu-hydro` [169], used in previous studies [26, 53, 170]. Both MUSIC and VISHNew solve the same hydrodynamic equations of motion [73] but with two different numerical schemes: VISHNew uses SHASTA [171] while MUSIC uses the Kurganov-Tadmor algorithm [52]. Despite differences in the numerical algorithms — amounting to approximations of spatial derivatives — for sufficiently smooth hydrodynamic fields the two codes should agree well.

Besides the different numerical schemes, VISHNew and MUSIC have different viscous current regulation schemes. The regulation scheme used in VISHNew is described in Ref. [74] while that used in MUSIC can be found in Ref. [15, 172]. For small to moderate values of η/s and ζ/s , neither of these schemes should regulate the viscous currents close to or inside the constant energy density (or temperature) switching hypersurface. Because our hydrodynamic model will explore moderate and large values of η/s and ζ/s , it is important to compare the hydrodynamic fields. For a fixed $\eta/s = 0.08$,

we run the same smooth initial conditions used for the ideal hydrodynamic comparison through free-streaming and either MUSIC or VISHNew with zero bulk viscosity and the conformal equation of state $\epsilon = 3p$. These are shown in Fig. 26.

At late times, there are differences in the shear stress π^{xy} near the dilute regions of the grid. These differences do not propagate into the region inside the particlization surface ($\epsilon \gtrsim 0.2 \text{ GeV/fm}^3$). We have also run the exact same event through viscous hydro with a fixed $\eta/s = 0.3$. The larger specific shear viscosity will incur stronger regulation. We find that the MUSIC scheme, while aggressive in low temperature regions, allows larger values of shear pressure inside the region $\epsilon > 0.2 \text{ GeV/fm}^3$.

As additional validation, we repeated the previous test with a QCD equation of state (as described in Section III B), again with a fixed specific shear viscosity $\eta/s = 0.08$ but this time with a temperature dependent specific bulk viscosity ($\zeta/s(T)$) from Ref. [53]. The bulk pressure, energy density and flow are shown in Figs. 27 and 28. Good agreement is found between the two codes.

In order to quantify the effects of any small differences that the hydrodynamics may have on our hadronic observables, we have evaluated the smooth Cooper-Frye integral over the switching surface generated by each hydrodynamics code. The hydrodynamic event used was the same event with bulk and shear pressures for which the hydrodynamic evolution was compared above. We used iS3D to perform the smooth Cooper-Frye integral over each surface, including bulk and shear Grad viscous corrections, and plotted the comparisons below for pions, kaons and protons. In general, the agreement in the spectra is very good. These are shown in Fig. 29. These differences of about 1% or less in the differential observables yield differences $\lesssim 1\%$ in the p_T and ϕ_p integrated observables.

a. Validation against cylindrically symmetric external solution

In this section, we compare the ideal hydrodynamic numerical solution obtained with MUSIC and VISHNew with a solution obtained with Mathematica [173]. Cylindrically symmetric initial conditions simplify considerably the equations of hydrodynamics [174], making them solvable in Mathematica [175]. Since the differential equation solver in Mathematica is adaptive and completely different from that of MUSIC and VISHNew, it provides complementary validation of their numerical solutions. We focus on a scenario where large gradients develop, to challenge the numerical algorithms of MUSIC and VISHNew. We use the initial temperature profile $T(\tau_0, r) = T_0 \exp(-r^2/\sigma^2)$ with $\tau_0 = 0.4 \text{ fm}/c$, $T_0 = 600 \text{ MeV}$ and $\sigma = 1 \text{ fm}$. We use a conformal equation of state. The result is shown in Fig. 30 for the temperature profile (a) and the flow rapidity u^x (b).

Overall, all three solutions agree well in regions of larger temperatures, which are the physically relevant regions of spacetime. Near the edge of the fireball, in regions of very low temperatures (small τ and large x), neither MUSIC and VISHNew can reconstruct the flow velocity well. This is a

well-know challenge of many numerical solution of relativistic hydrodynamic equation used in heavy ion physics: regions of very low temperatures are difficult to solve accurately and are susceptible to numerical instabilities. These issues at very low temperatures typically do not propagate to the higher-temperature regions and, as such, are not serious in practice.

2. SMASH

The use of SMASH as an afterburner for event-by-event studies of heavy-ion collisions is still fairly new. For this reason, we have made a comparison between UrQMD and SMASH as they relate to our transverse-momentum-integrated observables. We generated five thousand fluctuating initial conditions for Au-Au $\sqrt{s_{NN}} = 0.2 \text{ TeV}$ collisions with parameters fixed by the Maximum A Posteriori parameters found in [53] except for the initial energy density normalization, which was scaled to fit the multiplicities.

We allowed each initial condition to free-stream for the same time and then used these initial conditions for hydrodynamics in two different models:

1. SMASH: We matched the HotQCD lattice equation of state to the SMASH list of resonances (excluding the σ meson). Each initial condition was propagated through viscous hydrodynamics with this equation of state, followed by particlization using the Pratt-Torrieri-Bernhard viscous correction ansatz, followed by dynamics in SMASH.
2. UrQMD: We matched the HotQCD lattice equation of state to the list of resonances which can be propagated in UrQMD. Each initial condition was propagated through viscous hydrodynamics with this equation of state, followed by particlization using the Pratt-Torrieri-Bernhard viscous correction ansatz, followed by dynamics in UrQMD.

Finally we compared the observables predicted by the two models, shown in Fig. 31. For the observables we considered, we found very good agreement. In particular, heavier resonances have spectra that are more strongly influenced by the hadronic afterburner than lighter resonances and the agreement in the multiplicity and transverse momenta of the proton and Λ is strong.

We also ran hydrodynamics with a fixed equation of state matched to the SMASH hadron resonance gas particle content and then switched at the same temperature to UrQMD or SMASH. Because of the mismatch between the equation of state generated with the SMASH and the UrQMD resonance gases, there is a discrepancy at particlization in all of the thermodynamic variables. For example, at the same temperature, the energy density of the SMASH resonance gas and UrQMD's are different. This leads to a disagreement in observables. In particular, observables sensitive to the normalization of energy density, such as multiplicities and the transverse energy, showed a discrepancy at the level of approximately five percent. It is easy to understand that the energy density of the UrQMD resonance gas is a few percent smaller than SMASH's at the same temperature

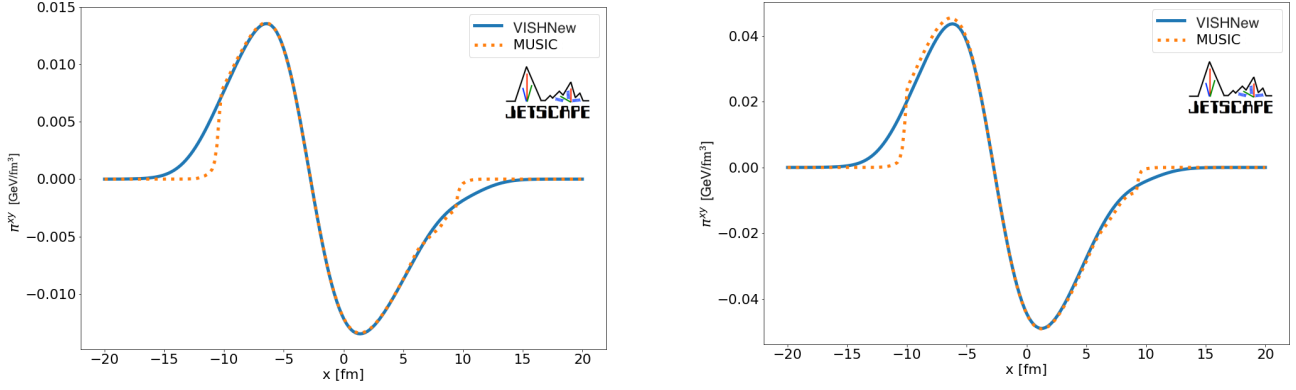


FIG. 26. The results of the hydrodynamic evolution of the shear stress for a smooth initial condition, just before freeze-out, for $\eta/s = 0.08$ (left) and $\eta/s = 0.3$ (right). The MUSIC regulation scheme allows larger inverse Reynolds numbers inside of the switching surface than the VISHNew scheme.




FIG. 27. The initial bulk pressure (left) and bulk pressure just before freeze-out (right), resulting from hydrodynamic evolution of a smooth initial condition. The specific shear viscosity was fixed $\eta/s = 0.08$, and specific bulk viscosity $(\zeta/s)(T)$ was given by [53] for this test.

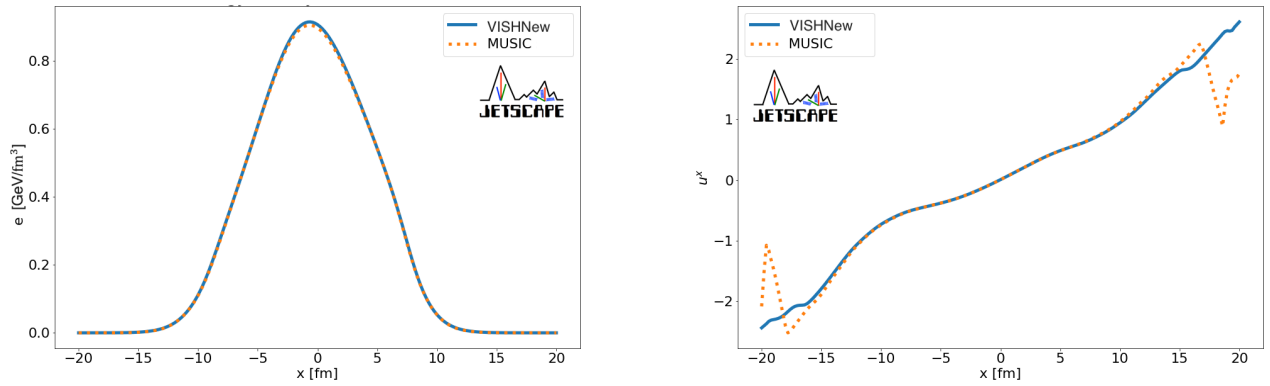


FIG. 28. The energy density (left) and flow (right) after hydrodynamic evolution of a smooth initial condition. The specific shear viscosity was fixed $\eta/s = 0.08$, and specific bulk viscosity $(\zeta/s)(T)$ was given by [53] for this test.

because of the different species and masses of hadrons. More details can be found in Appendix H6.

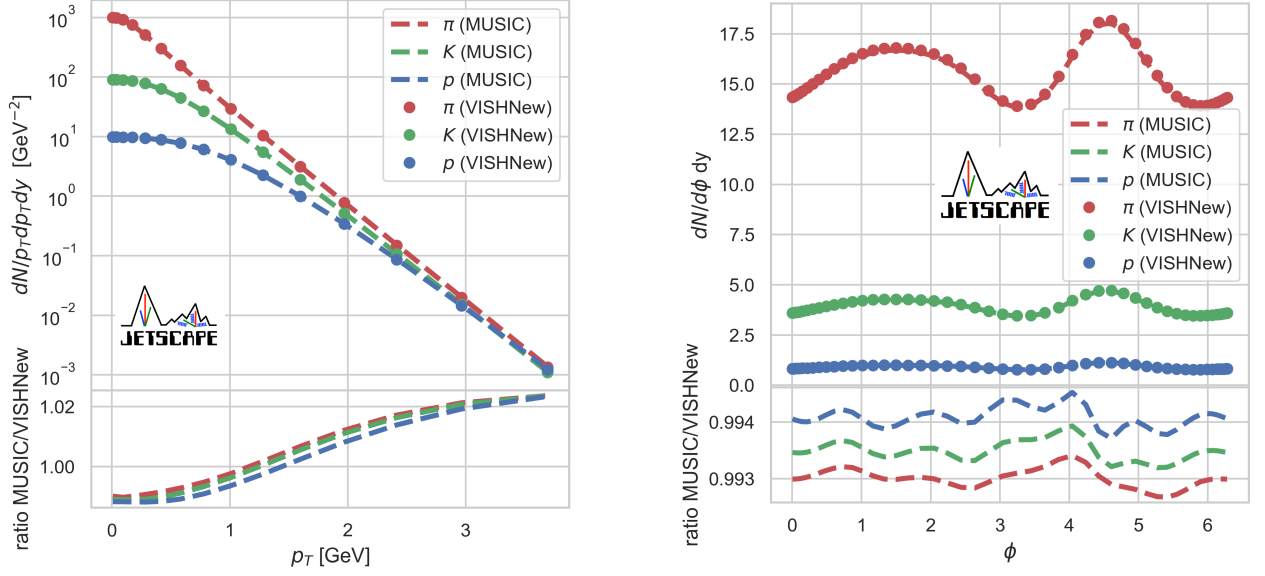


FIG. 29. Comparison of the transverse momentum p_T spectra (left) and azimuthal ϕ_p spectra (right) generated from the MUSIC and VISHNew freezeout surfaces. The freezeout surface was generated using the events compared above, with fixed $\eta/s = 0.08$, and specific bulk viscosity $(\zeta/s)(T)$ was given by [53].




FIG. 30. Temperature (left) and u^x profiles (right) for ideal hydrodynamics initialized with $T(\tau_0, r) = T_0 \exp(-r^2/\sigma^2)$ with $\tau_0 = 0.4 \text{ fm}/c$, $T_0 = 600 \text{ MeV}$ and $\sigma = 1 \text{ fm}$. Conformal equation of state. Result from new version of MUSIC and from VISHNew, compared with a solution obtained from Mathematica. The temperature contours in the left panel are (from left to right) for 10, 25, 50, 75, 100, 200 and 400 MeV, while the flow rapidity contours are (from top to bottom) for rapidities ranging from -4 to 4 by step of 1.

Given the novelty of using SMASH as an afterburner, we share for completeness the numerical parameters that we used with SMASH. These parameters, shown in Table V, gave sufficient accuracy without unreasonable loss of speed.

3. Comparison of JETSCAPE with hic-eventgen

In addition to validating of all the separate model components, we also have checked that the centrality-averaged observables predicted by our JETSCAPE model agree very well with a version of hic-eventgen, the event generator used in Ref. [53]. This was performed by restricting our parametriza-

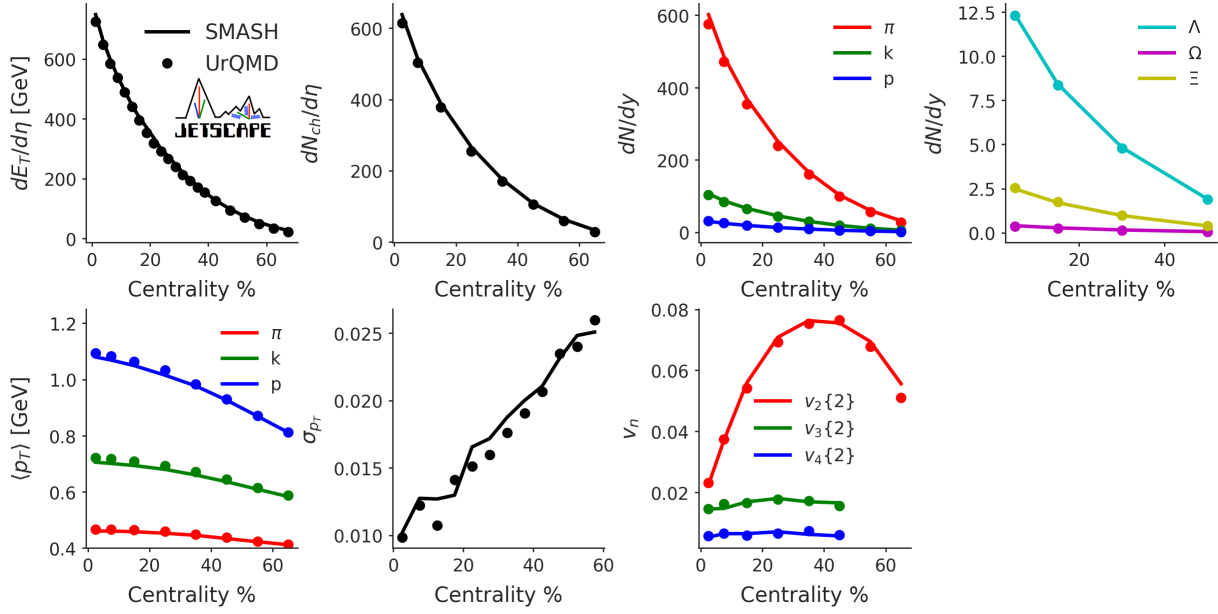


FIG. 31. Comparison of soft hadronic observables for Au-Au $\sqrt{s_{NN}} = 0.2$ TeV collisions using the SMASH (solid lines) or UrQMD (dots) afterburner.

Modus	Afterburner
Time_Step_Mode	Fixed
Delta_Time	1.0
End_Time	1000.0

TABLE V. SMASH parameters used event-by-event throughout this study.

σ	$dN_{ch}/d\eta$	$dE_T/d\eta$ [GeV]	dN_π/dy	$\langle p_T^\pi \rangle$ [GeV]
$m = 475$ MeV	615	777	569	0.54
$m = 800$ MeV	583	754	534	0.55
$m \rightarrow \infty$	579	743	531	0.54

TABLE VI. Sensitivity of observables to inclusion of the σ meson. The observables computed using the same cuts as the ALICE experiment.

tions to be the same as the Maximum A Posteriori parameters found in that study. The results of this comparison are shown in Fig. 32, in which we have averaged over five thousand fluctuating Pb-Pb $\sqrt{s_{NN}} = 2.76$ TeV collision events.

In general, we find excellent agreement between the two hybrid models. For this level of agreement, the σ meson had to be excluded from the `hic-eventgen` model; all resonances were also sampled on their mass-shell in `frzout` [176], the particle sampler in `hic-eventgen`. The equation of state used during the hydrodynamic evolution was constructed to match the hadron resonance gas used in `frzout` (excluding the σ meson).

4. The σ meson

The effects of including a σ meson resonance in our hadron resonance gas are studied using the `frzout` module [176], which is designed with the option to sample the σ resonance as a thermal resonance and perform its decay to pions. In particular, we compare three scenarios:

- Excluding the σ meson from sampling (labeled by $m \rightarrow \infty$).
- Sampling the σ meson with the PDG pole mass (~ 500

MeV) [177].

- Sampling the σ meson with the mass used in SMASH (~ 800 MeV).

The `frzout` module was used to sample particles from a hypersurface generated by the MUSIC simulation of a mid-central Pb-Pb event. The initial condition, free-streaming, and hydrodynamic transport parameters were set by the Maximum A Posteriori parameters given in [53]. The switching temperature was 151 MeV. Sampled particles are then propagated to UrQMD to perform hadronic rescatterings. Note that UrQMD does not have a σ meson: the effect of the σ meson is purely being tested at the level of the particlization, not in the afterburner. A total number of 100 over-samples were generated to increase the statistics. The results on charged-particle multiplicity, transverse energy, and pion multiplicity and mean transverse momentum are shown in Table VI.

Because the σ meson decays into pions, we see that the pion yield can differ by 7% for the lightest σ resonance. For the higher mass σ , the results are close to not sampling a σ meson. Additional differences would manifest if we included the effect of varying the σ mass in constructing the hadron gas equation of state as this would also have an effect on the hydrodynamic evolution. As has been explained in the main text, we chose

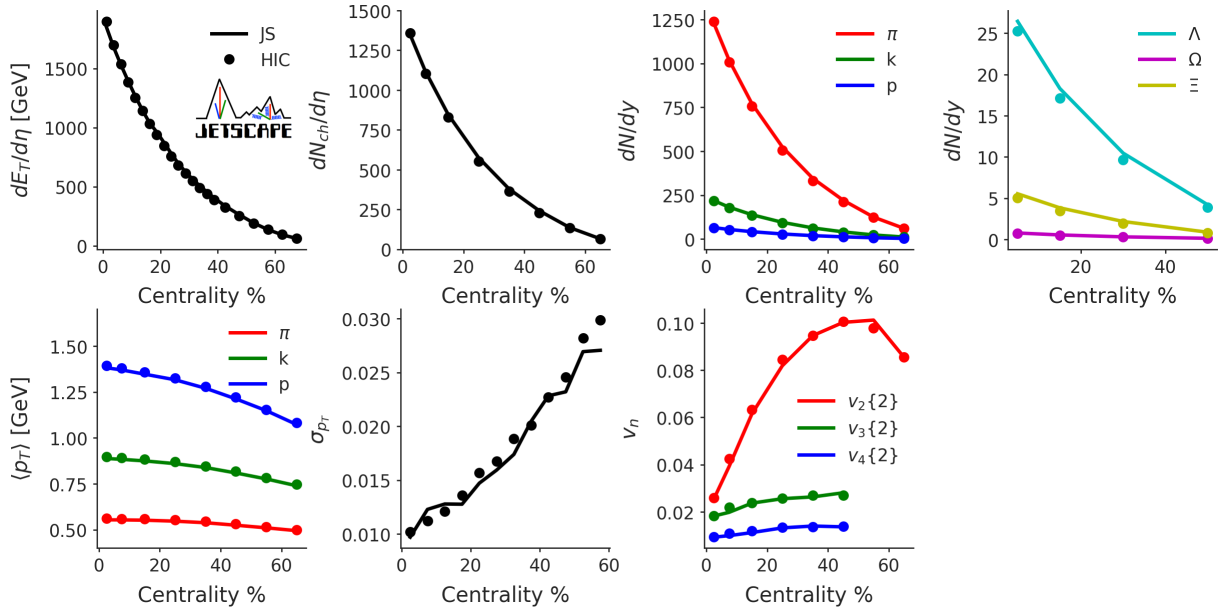


FIG. 32. Comparison of soft hadronic observables depending on whether one uses the JETSCAPE event generator (solid lines) or `hic-eventgen` (dots), averaged over five-thousand Pb-Pb $\sqrt{s_{NN}} = 2.76$ TeV events.

to omit the σ from the equation of state and particlization, following [125].

5. Sampling particles on mass-shell

In general, resonances occupying the hadron resonance gas should be described by spectral functions with a nonzero width. Previous Bayesian analyses have employed the sampling of thermal resonances off-shell [53, 170] in their extraction of medium properties. In this work, we chose to sample resonances on their pole-mass, neglecting this effect for now.

In what follows, we perform a simple test to quantify the effect of off-shell sampling of resonance masses for hadronic observables used in this work.²⁹ The results are shown in Fig. 33. The particles sampled off-shell are put on-shell after one collision in UrQMD. All particles occupying the hadron resonance gas are assigned Breit-Wigner spectral functions. The sampling was performed in both cases using the `frzout` module [176]; details about the hadron resonance gas pole-masses and widths can be found in Ref. [53].

We find a $\sim 5\%$ effect from the off-shell sampling in the hadronic observables shown in Fig. 33. We believe this effect is larger because of the σ meson present in this test: a light σ meson with pole mass $m \approx 475$ MeV and a very broad width $\Gamma \approx 500$ MeV, as used in Ref. [53]. As we saw in the previous section, the presence of the σ meson (sampled on the pole-mass) contributes significantly to the pion yield.




FIG. 33. Ratio of observables between sampling resonances with a non-zero width and sampling on the pole-mass. We find that the differences of less than 5% for the observables used in this study justify neglecting the non-zero width of resonances when performing parameter estimation.

6. QCD equations of state with different hadron resonance gases

The QCD equation of state used in hydrodynamic simulations of heavy ion collision matches a lattice calculation at high temperature ($T \gtrsim 120$ MeV) with a hadron resonance gas calculation at low temperature. In this work, we use the lattice calculations from Ref. [2]. As explained in Ref. [2], the trace anomaly evaluated from lattice is used to compute the pressure by integration of

$$\frac{p(T)}{T^4} = \frac{p_0(T)}{T_0^4} + \int_{T_0}^T \frac{dT'}{T'} \frac{\Theta^{\mu\mu}(T')}{T'^4}. \quad (\text{H1})$$

Energy density and entropy density then follow. The integration constant for the pressure is obtained from a hadron resonance gas calculation at $T = 130$ MeV.

Reference [53] followed a related but different approach: to ensure energy-momentum conservation at particlization, the

²⁹ This test used 4000 Pb-Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV events within a narrow impact parameter range $9.88 \text{ fm} < b < 11.02 \text{ fm}$.

lattice QCD trace anomaly is matched to a hadron resonance gas in a $[T_a, T_b]$ temperature range. The trace anomaly below T_a is that of the hadron resonance gas; the trace anomaly between T_a and T_b is an interpolation between the resonance gas and the lattice QCD trace anomaly. Above T_b , the trace anomaly is that of the lattice QCD. Using this new trace anomaly, which differs from that of the lattice below T_b , the pressure is computed by integration using $p_0(T_0 = 50\text{MeV})$ as reference; energy density and entropy density are then calculated.

We illustrate first the differences between the lattice pressure, and the pressure obtained with the above matching. If the temperature is below the matching point T_a , the pressure from the lattice case is given by

$$\frac{p_L(T)}{T^4} = \frac{p_{L,0}}{T_0^4} + \int_{T_0}^T \frac{dT'}{T'} \frac{\Theta_L^{\mu\mu}(T')}{T'^4}. \quad (\text{H2})$$

where the integration constant $p_{L,0}$ is the only input from the hadron resonance gas that enters in the definition of the pressure.

In the ‘‘matched’’ equation of state, however, the entire thermodynamics is determined by the hadron resonance gas below the lower matching temperature T_a :

$$\frac{p_M(T)}{T^4} = \frac{p_{M,0}}{T_0^4} + \int_{T_0}^T \frac{dT'}{T'} \frac{\Theta_{HRG}^{\mu\mu}(T')}{T'^4}. \quad (\text{H3})$$

There is no information from the lattice calculations entering in Eq. (H3) if $T < T_a$. This example makes it clear that any mismatch between the trace anomaly of lattice calculations and that of the hadron resonance gas results in a difference in the equation of state. This is of course the case even if the exact same hadron resonance gas are used to fix $p_{L,0} - p_{L,0} = p_{M,0}$ — which is arguably never the case. These uncertainties are difficult to eliminate: any mismatch between the hadron resonance gas and the lattice calculation would result in a discontinuity at particlization.

Evidently, even with the same matching procedure between the hadron resonance gas and the lattice calculation, the exact content of the hadron resonance gas is important. In the present case, we are interested in two configurations: one used the particle content from SMASH, while the other uses UrQMD’s. Both are matched to HotQCD’s lattice calculation as described above. The equation of state with the SMASH hadron resonance gas is the one that has been used to perform parameter estimation in this study. We compare the two equations of state in Figs. 34 and 35. The differences between the two equations of state amount to up to 8%.

In Appendix H2, we compared the predictions of two hybrid models, one model using SMASH as afterburner and the other using UrQMD. To obtain such a level of agreement in

the observables, it was necessary to use, in the hydrodynamics, equations of state that matched consistently the chosen hadronic transport afterburner. This is consistent with what we see in Fig. 34: inside the window of particlization temperature, the differences between the equations of state can be larger than $\sim 5\%$, and can undeniably produce noticeably different hadronic observables.

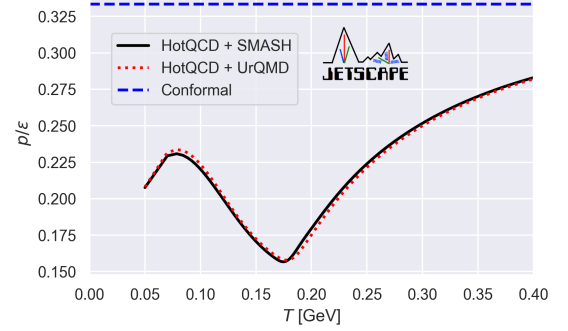


FIG. 34. The equation of state used throughout this work for parameter estimation ‘HotQCD + SMASH’ is shown as well as a different equation of state that has been matched to the list of resonances propagated in UrQMD. The conformal equation of state is included as a visual reference.




FIG. 35. Ratio of SMASH / UrQMD at the same temperature for three thermodynamic quantities: energy density (red), equilibrium pressure (green), and entropy density (blue). Each equation of state is constructed by matching the l QCD equation of state to a hadron resonance gas matching the list/masses of particles for each code. We see that the disagreement is largest near the region of the switching temperature.

We note that this work uses a fixed equation of state which does not parametrize any potential theoretical uncertainties. See Ref. [76] for a recent study which includes uncertainty in the lattice-matched equation of state.

[1] S. Borsanyi, Z. Fodor, C. Hoelbling, S. D. Katz, S. Krieg, and K. K. Szabo, *Phys. Lett.* **B730**, 99 (2014), arXiv:1309.5258

[hep-lat].

- [2] A. Bazavov *et al.* (HotQCD), *Phys. Rev.* **D90**, 094503 (2014), [arXiv:1407.6387 \[hep-lat\]](#).
- [3] C. Shen, in *28th International Conference on Ultrarelativistic Nucleus-Nucleus Collisions* (2020) [arXiv:2001.11858 \[nucl-th\]](#).
- [4] A. Bazavov, F. Karsch, S. Mukherjee, and P. Petreczky (USQCD), *Eur. Phys. J. A* **55**, 194 (2019), [arXiv:1904.09951 \[hep-lat\]](#).
- [5] E. Lu and G. D. Moore, *Phys. Rev. C* **83**, 044901 (2011), [arXiv:1102.0017 \[hep-ph\]](#).
- [6] J. B. Rose, J. Torres-Rincon, A. Schäfer, D. Oliinychenko, and H. Petersen, *Phys. Rev. C* **97**, 055204 (2018), [arXiv:1709.03826 \[nucl-th\]](#).
- [7] J.-B. Rose, J. Torres-Rincon, and H. Elfner, (2020), [arXiv:2005.03647 \[hep-ph\]](#).
- [8] F. Karsch, D. Kharzeev, and K. Tuchin, *Phys. Lett. B* **663**, 217 (2008), [arXiv:0711.0914 \[hep-ph\]](#).
- [9] J. Noronha-Hostler, J. Noronha, and C. Greiner, *Phys. Rev. Lett.* **103**, 172302 (2009), [arXiv:0811.1571 \[nucl-th\]](#).
- [10] P. B. Arnold, C. Dogan, and G. D. Moore, *Phys. Rev. D* **74**, 085021 (2006), [arXiv:hep-ph/0608012](#).
- [11] J. Ghiglieri, G. D. Moore, and D. Teaney, *JHEP* **03**, 179 (2018), [arXiv:1802.09535 \[hep-ph\]](#).
- [12] U. Heinz and R. Snellings, *Ann. Rev. Nucl. Part. Sci.* **63**, 123 (2013), [arXiv:1301.2826 \[nucl-th\]](#).
- [13] C. Gale, S. Jeon, and B. Schenke, *Int. J. Mod. Phys. A* **28**, 1340011 (2013), [arXiv:1301.5893 \[nucl-th\]](#).
- [14] R. Derradi de Souza, T. Koide, and T. Kodama, *Prog. Part. Nucl. Phys.* **86**, 35 (2016), [arXiv:1506.03863 \[nucl-th\]](#).
- [15] B. Schenke, C. Shen, and P. Tribedy, (2020), [arXiv:2005.14682 \[nucl-th\]](#).
- [16] P. Romatschke and U. Romatschke, *Phys. Rev. Lett.* **99**, 172301 (2007), [arXiv:0706.1522 \[nucl-th\]](#).
- [17] H. Song and U. Heinz, *Phys. Lett. B* **658**, 279 (2008), [arXiv:0709.0742 \[nucl-th\]](#).
- [18] H. Song and U. Heinz, in *Strangeness in Quark Matter 2008. Proceedings, International Conference, SQM 2008, Beijing, P.R. China, October 5-10, 2008*, Vol. G36 (2009) p. 064033, [arXiv:0812.4274 \[nucl-th\]](#).
- [19] G. S. Denicol, T. Kodama, T. Koide, and P. Mota, *Phys. Rev. C* **80**, 064901 (2009), [arXiv:0903.3595 \[hep-ph\]](#).
- [20] “Hot and Dense QCD White Paper,” http://www.bnl.gov/npp/docs/Bass_RHI_WP_final.pdf (2012).
- [21] Y. Akiba *et al.*, (2015), [arXiv:1502.02730 \[nucl-ex\]](#).
- [22] H. Petersen, C. Coleman-Smith, S. A. Bass, and R. Wolpert, *J. Phys. G* **38**, 045102 (2011), [arXiv:1012.4629 \[nucl-th\]](#).
- [23] J. Novak, K. Novak, S. Pratt, J. Vredevoogd, C. Coleman-Smith, and R. Wolpert, *Phys. Rev. C* **89**, 034917 (2014), [arXiv:1303.5769 \[nucl-th\]](#).
- [24] E. Sangaline and S. Pratt, *Phys. Rev. C* **93**, 024908 (2016), [arXiv:1508.07017 \[nucl-th\]](#).
- [25] J. E. Bernhard, P. W. Marcy, C. E. Coleman-Smith, S. Huzar-bazar, R. L. Wolpert, and S. A. Bass, *Phys. Rev. C* **91**, 054910 (2015), [arXiv:1502.00339 \[nucl-th\]](#).
- [26] J. E. Bernhard, J. S. Moreland, S. A. Bass, J. Liu, and U. Heinz, *Phys. Rev. C* **94**, 024907 (2016), [arXiv:1605.03954 \[nucl-th\]](#).
- [27] J. E. Bernhard, J. S. Moreland, and S. A. Bass, *Nature Phys.* **15**, 1113 (2019).
- [28] D. Everett *et al.* (JETSCAPE), (2020), [arXiv:2011.01430 \[hep-ph\]](#).
- [29] H. Song, S. A. Bass, U. Heinz, T. Hirano, and C. Shen, *Phys. Rev. Lett.* **106**, 192301 (2011), [Erratum: *Phys.Rev.Lett.* 109, 139904 (2012)], [arXiv:1011.2783 \[nucl-th\]](#).
- [30] G. Nijs, W. van der Schee, U. Gürsoy, and R. Snellings, (2020), [arXiv:2010.15130 \[nucl-th\]](#).
- [31] G. Nijs, W. van der Schee, U. Gürsoy, and R. Snellings, (2020), [arXiv:2010.15134 \[nucl-th\]](#).
- [32] J. H. Putschke *et al.* (JETSCAPE), (2019), [arXiv:1903.07706 \[nucl-th\]](#).
- [33] D. S. Sivia and J. Skilling, *Data Analysis - A Bayesian Tutorial*, 2nd ed., Oxford Science Publications (Oxford University Press, 2006).
- [34] S. Bass and A. Dumitru, *Phys. Rev. C* **61**, 064909 (2000), [arXiv:nucl-th/0001033](#).
- [35] C. Nonaka and S. A. Bass, *Phys. Rev. C* **75**, 014902 (2007), [arXiv:nucl-th/0607018 \[nucl-th\]](#).
- [36] T. Hirano, U. Heinz, D. Kharzeev, R. Lacey, and Y. Nara, *Phys. Rev. C* **77**, 044909 (2008), [arXiv:0710.5795 \[nucl-th\]](#).
- [37] H. Petersen, J. Steinheimer, G. Burau, M. Bleicher, and H. Stocker, *Phys. Rev. C* **78**, 044901 (2008), [arXiv:0806.1695 \[nucl-th\]](#).
- [38] H. Song, S. A. Bass, and U. Heinz, *Phys. Rev. C* **83**, 024912 (2011), [arXiv:1012.0555 \[nucl-th\]](#).
- [39] U. Heinz, C. Shen, and H. Song, *AIP Conf. Proc.* **1441**, 766 (2012), [arXiv:1108.5323 \[nucl-th\]](#).
- [40] H. Song, S. Bass, and U. Heinz, *Phys. Rev. C* **89**, 034919 (2014), [arXiv:1311.0157 \[nucl-th\]](#).
- [41] X. Zhu, F. Meng, H. Song, and Y.-X. Liu, *Phys. Rev. C* **91**, 034904 (2015), [arXiv:1501.03286 \[nucl-th\]](#).
- [42] S. Ryu, J.-F. Paquet, C. Shen, G. Denicol, B. Schenke, S. Jeon, and C. Gale, *Phys. Rev. C* **97**, 034910 (2018), [arXiv:1704.04216 \[nucl-th\]](#).
- [43] M. Vardanyan, R. Trotta, and J. Silk, *Monthly Notices of the Royal Astronomical Society: Letters* **413**, L91–L95 (2011).
- [44] J. S. Moreland, J. E. Bernhard, and S. A. Bass, *Phys. Rev. C* **92**, 011901 (2015), [arXiv:1412.4708 \[nucl-th\]](#).
- [45] <https://github.com/Duke-QCD/trento.git>.
- [46] J. Liu, C. Shen, and U. Heinz, *Phys. Rev. C* **91**, 064906 (2015), [Erratum: *Phys. Rev.C*92,no.4,049904(2015)], [arXiv:1504.02160 \[nucl-th\]](#).
- [47] W. Broniowski, W. Florkowski, M. Chojnacki, and A. Kisiel, *Phys. Rev. C* **80**, 034902 (2009), [arXiv:0812.3393 \[nucl-th\]](#).
- [48] <https://github.com/derekeverett/freestream-milne>.
- [49] B. Schenke, S. Jeon, and C. Gale, *Phys. Rev. C* **82**, 014903 (2010), [arXiv:1004.1408 \[hep-ph\]](#).
- [50] B. Schenke, S. Jeon, and C. Gale, *Phys. Rev. Lett.* **106**, 042301 (2011), [arXiv:1009.3244 \[hep-ph\]](#).
- [51] J.-F. Paquet, C. Shen, G. S. Denicol, M. Luzum, B. Schenke, S. Jeon, and C. Gale, *Phys. Rev. C* **93**, 044906 (2016), [arXiv:1509.06738 \[hep-ph\]](#).
- [52] A. Kurganov and E. Tadmor, *Journal of Computational Physics* **160**, 241 (2000).
- [53] J. E. Bernhard, *Bayesian parameter estimation for relativistic heavy-ion collisions*, Ph.D. thesis, Duke U. (2018-04-19), [arXiv:1804.06469 \[nucl-th\]](#).
- [54] https://github.com/j-f-paquet/eos_maker.
- [55] M. McNelis, D. Everett, and U. Heinz, *Comput. Phys. Commun.* **258**, 107604 (2021), [arXiv:1912.08271 \[nucl-th\]](#).
- [56] <https://github.com/derekeverett/iS3D>.
- [57] J. Weil *et al.*, *Phys. Rev. C* **94**, 054905 (2016), [arXiv:1606.06642 \[nucl-th\]](#).
- [58] <https://github.com/smash-transport/smash>.
- [59] W. Florkowski, M. P. Heller, and M. Spalinski, *Rept. Prog. Phys.* **81**, 046001 (2018), [arXiv:1707.02282 \[hep-ph\]](#).
- [60] J. Vredevoogd and S. Pratt, *Phys. Rev. C* **79**, 044915 (2009), [arXiv:0810.4325 \[nucl-th\]](#).

- [61] L. Keegan, A. Kurkela, A. Mazeliauskas, and D. Teaney, *JHEP* **08**, 171 (2016), arXiv:1605.04287 [hep-ph].
- [62] A. Kurkela, A. Mazeliauskas, J.-F. Paquet, S. Schlichting, and D. Teaney, *Phys. Rev.* **C99**, 034910 (2019), arXiv:1805.00961 [hep-ph].
- [63] A. Kurkela, A. Mazeliauskas, J.-F. Paquet, S. Schlichting, and D. Teaney, *Phys. Rev. Lett.* **122**, 122302 (2019), arXiv:1805.01604 [hep-ph].
- [64] S. Schlichting and D. Teaney, *Ann. Rev. Nucl. Part. Sci.* **69**, 447 (2019), arXiv:1908.02113 [nucl-th].
- [65] A. Kurkela, W. van der Schee, U. A. Wiedemann, and B. Wu, *Phys. Rev. Lett.* **124**, 102301 (2020), arXiv:1907.08101 [hep-ph].
- [66] W. Ke, J. S. Moreland, J. E. Bernhard, and S. A. Bass, *Phys. Rev.* **C96**, 044912 (2017), arXiv:1610.08490 [nucl-th].
- [67] J. S. Moreland, J. E. Bernhard, and S. A. Bass, *Phys. Rev. C* **101**, 024911 (2020), arXiv:1808.02106 [nucl-th].
- [68] B. Schenke, P. Tribedy, and R. Venugopalan, *Phys. Rev. Lett.* **108**, 252301 (2012), arXiv:1202.6646 [nucl-th].
- [69] C. Shen and S. Alzhrani, (2020), arXiv:2003.05852 [nucl-th].
- [70] M. L. Miller, K. Reygers, S. J. Sanders, and P. Steinberg, *Ann. Rev. Nucl. Part. Sci.* **57**, 205 (2007), arXiv:nucl-ex/0701025 [nucl-ex].
- [71] D. Oliinychenko and H. Petersen, *Phys. Rev. C* **93**, 034905 (2016), arXiv:1508.04378 [nucl-th].
- [72] G. Başar and D. Teaney, *Phys. Rev. C* **90**, 054903 (2014), arXiv:1312.6770 [nucl-th].
- [73] G. S. Denicol, E. Molnár, H. Niemi, and D. H. Rischke, *Eur. Phys. J.* **A48**, 170 (2012), arXiv:1206.1554 [nucl-th].
- [74] C. Shen, Z. Qiu, H. Song, J. Bernhard, S. Bass, and U. Heinz, *Comput. Phys. Commun.* **199**, 61 (2016), arXiv:1409.8164 [nucl-th].
- [75] <http://www.physics.mcgill.ca/music/>.
- [76] J. Auvinen, K. J. Eskola, P. Huovinen, H. Niemi, R. Paatelainen, and P. Petreczky, (2020), arXiv:2006.12499 [nucl-th].
- [77] J. Liu, *Pre-equilibrium evolution effects on relativistic heavy-ion collision observables*, Ph.D. thesis, Ohio State University (2015).
- [78] B. Schenke, C. Shen, and P. Tribedy, (2019), arXiv:1908.06212 [nucl-th].
- [79] G. Denicol, *J. Phys. G* **41**, 124004 (2014).
- [80] G. S. Denicol, S. Jeon, and C. Gale, *Phys. Rev.* **C90**, 024912 (2014), arXiv:1403.0962 [nucl-th].
- [81] H. Niemi, K. Eskola, and R. Paatelainen, *Phys. Rev. C* **93**, 024907 (2016), arXiv:1505.02677 [hep-ph].
- [82] L. P. Csernai, J. Kapusta, and L. D. McLerran, *Phys. Rev. Lett.* **97**, 152303 (2006), arXiv:nucl-th/0604032.
- [83] D. Kharzeev and K. Tuchin, *JHEP* **09**, 093 (2008), arXiv:0705.4280 [hep-ph].
- [84] R. Baier, P. Romatschke, D. T. Son, A. O. Starinets, and M. A. Stephanov, *JHEP* **04**, 100 (2008), arXiv:0712.2451 [hep-th].
- [85] S. Bhattacharyya, V. E. Hubeny, S. Minwalla, and M. Rangamani, *JHEP* **02**, 045 (2008), arXiv:0712.2456 [hep-th].
- [86] W. Florkowski, A. Jaiswal, E. Maksymiuk, R. Ryblewski, and M. Strickland, *Phys. Rev.* **C91**, 054907 (2015), arXiv:1503.03226 [nucl-th].
- [87] A. Czajka, S. Hauksson, C. Shen, S. Jeon, and C. Gale, *Phys. Rev.* **C97**, 044914 (2018).
- [88] J. Ghiglieri, G. D. Moore, and D. Teaney, *Phys. Rev. Lett.* **121**, 052302 (2018), arXiv:1805.02663 [hep-ph].
- [89] S. Pu, T. Koide, and D. H. Rischke, *Phys. Rev.* **D81**, 114039 (2010), arXiv:0907.3906 [hep-ph].
- [90] H. Song and U. Heinz, *Phys. Rev. C* **78**, 024902 (2008), arXiv:0805.1756 [nucl-th].
- [91] T. Schaefer, *Ann. Rev. Nucl. Part. Sci.* **64**, 125 (2014), arXiv:1403.0653 [hep-ph].
- [92] F. Cooper and G. Frye, *Phys. Rev. D* **10**, 186 (1974).
- [93] F. Cooper, G. Frye, and E. Schonberg, *Phys. Rev.* **D11**, 192 (1975).
- [94] P. Huovinen and H. Petersen, *Eur. Phys. J.* **A48**, 171 (2012), arXiv:1206.3371 [nucl-th].
- [95] K. Dusling, G. D. Moore, and D. Teaney, *Phys. Rev.* **C81**, 034907 (2010), arXiv:0909.0754 [nucl-th].
- [96] D. Molnar and Z. Wolff, *Phys. Rev. C* **95**, 024903 (2017), arXiv:1404.7850 [nucl-th].
- [97] M. Damodaran, D. Molnar, G. G. Barnaföldi, D. Berényi, and M. Ferenc Nagy-Egri, (2017), arXiv:1707.00793 [nucl-th].
- [98] M. Damodaran, D. Molnar, G. G. Barnaföldi, D. Berényi, and M. F. Nagy-Egri, *Phys. Rev. C* **102**, 014907 (2020).
- [99] U. Heinz, *Nucl. Phys. A* **661**, 140 (1999), arXiv:nucl-th/9907060.
- [100] A. Andronic, P. Braun-Munzinger, and J. Stachel, *Nucl. Phys. A* **772**, 167 (2006), arXiv:nucl-th/0511071.
- [101] A. Andronic, P. Braun-Munzinger, K. Redlich, and J. Stachel, *Nature* **561**, 321 (2018), arXiv:1710.09425 [nucl-th].
- [102] H. Grad, *Commun. Pure Appl. Math* **2**, 331 (1949).
- [103] W. Israel, *Annals Phys.* **100**, 310 (1976).
- [104] W. Israel and J. M. Stewart, *Annals Phys.* **118**, 341 (1979).
- [105] D. Teaney, *Phys. Rev.* **C68**, 034913 (2003), arXiv:nucl-th/0301099 [nucl-th].
- [106] A. Monnai and T. Hirano, *Phys. Rev.* **C80**, 054906 (2009), arXiv:0903.4436 [nucl-th].
- [107] K. Dusling and T. Schäfer, *Phys. Rev.* **C85**, 044909 (2012), arXiv:1109.5181 [hep-ph].
- [108] G. S. Denicol, H. Niemi, E. Molnar, and D. H. Rischke, *Phys. Rev.* **D85**, 114047 (2012), [Erratum: *Phys. Rev. D*91,no.3,039902(2015)], arXiv:1202.4551 [nucl-th].
- [109] S. Chapman, T. G. Cowling, and D. Burnett, *The mathematical theory of non-uniform gases: an account of the kinetic theory of viscosity, thermal conduction and diffusion in gases* (Cambridge university press, 1990).
- [110] J. Anderson and H. Witting, *Physica* **74**, 466 (1974).
- [111] A. Jaiswal, R. Ryblewski, and M. Strickland, *Phys. Rev.* **C90**, 044908 (2014), arXiv:1407.7231 [hep-ph].
- [112] S. Pratt and G. Torrieri, *Phys. Rev.* **C82**, 044901 (2010), arXiv:1003.0413 [nucl-th].
- [113] P. L. Bhatnagar, E. P. Gross, and M. Krook, *Phys. Rev.* **94**, 511 (1954).
- [114] P. Romatschke and M. Strickland, *Phys. Rev.* **D68**, 036004 (2003), arXiv:hep-ph/0304092 [hep-ph].
- [115] P. Romatschke and M. Strickland, *Phys. Rev.* **D70**, 116006 (2004), arXiv:hep-ph/0406188 [hep-ph].
- [116] M. Martinez, R. Ryblewski, and M. Strickland, *Phys. Rev.* **C85**, 064913 (2012), arXiv:1204.1473 [nucl-th].
- [117] W. Florkowski, R. Ryblewski, and M. Strickland, *Phys. Rev.* **C88**, 024903 (2013), arXiv:1305.7234 [nucl-th].
- [118] W. Florkowski, R. Ryblewski, M. Strickland, and L. Tinti, *Phys. Rev.* **C89**, 054909 (2014), arXiv:1403.1223 [hep-ph].
- [119] L. Tinti, *Phys. Rev.* **C94**, 044902 (2016), arXiv:1506.07164 [hep-ph].
- [120] E. Molnar, H. Niemi, and D. H. Rischke, *Phys. Rev.* **D93**, 114025 (2016), arXiv:1602.00573 [nucl-th].
- [121] L. Tinti, G. Vujanovic, J. Noronha, and U. Heinz, *Proceedings, 27th International Conference on Ultrarelativistic Nucleus-Nucleus Collisions (Quark Matter 2018): Venice, Italy, May 14-19, 2018*, *Nucl. Phys. A* **982**, 919 (2019), arXiv:1808.06212 [nucl-th].

- [122] L. Tinti, G. Vujanovic, J. Noronha, and U. Heinz, *Phys. Rev.* **D99**, 016009 (2019), arXiv:1808.06436 [nucl-th].
- [123] T. Hirano and K. Tsuda, *Phys. Rev. C* **66**, 054905 (2002), arXiv:nucl-th/0205043.
- [124] J. Steinheimer, J. Aichelin, and M. Bleicher, *Phys. Rev. Lett.* **110**, 042501 (2013), arXiv:1203.5302 [nucl-th].
- [125] W. Broniowski, F. Giacosa, and V. Begun, *Phys. Rev.* **C92**, 034905 (2015), arXiv:1506.01260 [nucl-th].
- [126] E. T. Jaynes, *Phys. Rev.* **106**, 620 (1957).
- [127] F. Gelis, *Acta Phys. Polon.* **B 45**, 2257 (2014).
- [128] S. Bass, A. Dumitru, M. Bleicher, L. Bravina, E. Zabrodin, H. Stoecker, and W. Greiner, *Phys. Rev. C* **60**, 021902 (1999), arXiv:nucl-th/9902062.
- [129] J. Loeppky, J. Sacks, and W. Welch, *Technometrics* **51**, 366 (2009).
- [130] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, *Journal of Machine Learning Research* **12**, 2825 (2011).
- [131] J. Brynjarsdóttir and A. O'Hagan, *Inverse Problems* **30** (2014), 10.1088/0266-5611/30/11/114007.
- [132] D. W. Hogg and D. Foreman-Mackey, *The Astrophysical Journal Supplement Series* **236**, 11 (2018).
- [133] S. Brooks, A. Gelman, G. Jones, and X. Meng, *Handbook of Markov Chain Monte Carlo*, Chapman & Hall/CRC Handbooks of Modern Statistical Methods (CRC Press, 2011).
- [134] W. D. Voursden, W. M. Farr, and I. Mandel, *Monthly Notices of the Royal Astronomical Society* **455**, 1919 (2015).
- [135] C. Shen, U. Heinz, J.-F. Paquet, and C. Gale, *Phys. Rev. C* **89**, 044910 (2014), arXiv:1308.2440 [nucl-th].
- [136] K. Chaloner and I. Verdinelli, *Statist. Sci.* **10**, 273 (1995).
- [137] K. Aamodt *et al.* (ALICE), *Phys. Rev. Lett.* **106**, 032301 (2011), arXiv:1012.1657 [nucl-ex].
- [138] J. Adam *et al.* (ALICE), *Phys. Rev.* **C94**, 034903 (2016), arXiv:1603.04775 [nucl-ex].
- [139] B. Abelev *et al.* (ALICE), *Phys. Rev.* **C88**, 044910 (2013), arXiv:1303.0737 [hep-ex].
- [140] K. Aamodt *et al.* (ALICE), *Phys. Rev. Lett.* **107**, 032301 (2011), arXiv:1105.3865 [nucl-ex].
- [141] B. B. Abelev *et al.* (ALICE), *Eur. Phys. J.* **C74**, 3077 (2014), arXiv:1407.5530 [nucl-ex].
- [142] B. I. Abelev *et al.* (STAR), *Phys. Rev.* **C79**, 034909 (2009), arXiv:0808.2041 [nucl-ex].
- [143] J. Adams *et al.* (STAR), *Phys. Rev.* **C72**, 014904 (2005), arXiv:nucl-ex/0409033 [nucl-ex].
- [144] L. Adamczyk *et al.* (STAR), *Phys. Rev.* **C88**, 014904 (2013), arXiv:1301.2187 [nucl-ex].
- [145] S. Adler *et al.* (PHENIX), *Phys. Rev. C* **69**, 034909 (2004), arXiv:nucl-ex/0307022.
- [146] S. Pratt, E. Sangaline, P. Sorensen, and H. Wang, *Phys. Rev. Lett.* **114**, 202301 (2015), arXiv:1501.04042 [nucl-th].
- [147] Z. Wolff and D. Molnar, *Phys. Rev. C* **96**, 044909 (2017), arXiv:1611.09185 [nucl-th].
- [148] T. Lappi, *Phys. Lett. B* **643**, 11 (2006), arXiv:hep-ph/0606207.
- [149] C. Gale, S. Jeon, B. Schenke, P. Tribedy, and R. Venugopalan, *Phys. Rev. Lett.* **110**, 012302 (2013), arXiv:1209.6330 [nucl-th].
- [150] S. McDonald, C. Shen, F. Fillion-Gourdeau, S. Jeon, and C. Gale, *Phys. Rev. C* **95**, 064913 (2017), arXiv:1609.02958 [hep-ph].
- [151] T. Nunes da Silva, D. Chinellato, M. Hippert, W. Serenone, J. Takahashi, G. S. Denicol, M. Luzum, and J. Noronha, (2020), arXiv:2006.02324 [nucl-th].
- [152] M. Luzum and P. Romatschke, *Phys. Rev.* **C78**, 034915 (2008), [Erratum: *Phys. Rev. C* **79**, 039903 (2009)], arXiv:0804.4015 [nucl-th].
- [153] H. Song, *Causal Viscous Hydrodynamics for Relativistic Heavy Ion Collisions*, Ph.D. thesis, Ohio State U. (2009), arXiv:0908.3656 [nucl-th].
- [154] D. Hamby, *Environ Monit Assess* **32**, 135 (1994).
- [155] R. Trotta, *Contemp. Phys.* **49**, 71 (2008), arXiv:0803.4089 [astro-ph].
- [156] <https://github.com/willvousedn/ptemcee>.
- [157] P. Marshall, N. Rajguru, and A. Slosar, *Physical Review D* **73** (2006), 10.1103/physrevd.73.067302.
- [158] E. Schnedermann, J. Sollfrank, and U. Heinz, *Phys. Rev. C* **48**, 2462 (1993), arXiv:nucl-th/9307020.
- [159] U. Heinz, in *2nd CERN-CLAF School of High Energy Physics* (2004) pp. 165–238, arXiv:hep-ph/0407360.
- [160] <http://jetscape.org/sims-widget>.
- [161] U. A. Wiedemann, B. Tomasik, and U. Heinz, *Nucl. Phys. A* **638**, 475C (1998), arXiv:nucl-th/9801017.
- [162] G. Vujanovic, G. S. Denicol, M. Luzum, S. Jeon, and C. Gale, *Phys. Rev.* **C98**, 014902 (2018), arXiv:1702.02941 [nucl-th].
- [163] C. Gale, S. Jeon, S. McDonald, J.-F. Paquet, and C. Shen, *Nucl. Phys. A* **982**, 767 (2019), arXiv:1807.09326 [nucl-th].
- [164] S. Hauksson, S. Jeon, and C. Gale, in *28th International Conference on Ultrarelativistic Nucleus-Nucleus Collisions* (2020) arXiv:2001.10046 [hep-ph].
- [165] C. Gale, J.-F. Paquet, B. Schenke, and C. Shen, in *28th International Conference on Ultrarelativistic Nucleus-Nucleus Collisions* (2020) arXiv:2002.05191 [hep-ph].
- [166] J. S. Moreland and R. A. Soltz, *Phys. Rev. C* **93**, 044913 (2016), arXiv:1512.02189 [nucl-th].
- [167] <http://www.tacc.utexas.edu>.
- [168] “Ohio Supercomputer Center,” <http://osc.edu/ark:/19495/f5s1ph73> (1987).
- [169] <https://github.com/jbernhard/osu-hydro>.
- [170] J. S. Moreland, *Initial conditions of bulk matter in ultrarelativistic nuclear collisions*, Ph.D. thesis, Duke U. (2019), arXiv:1904.08290 [nucl-th].
- [171] J. P. Boris and D. L. Book, *J. Comput. Phys.* **11**, 38 (1973).
- [172] G. S. Denicol, C. Gale, S. Jeon, A. Monnai, B. Schenke, and C. Shen, *Phys. Rev.* **C98**, 034916 (2018), arXiv:1804.10557 [nucl-th].
- [173] Wolfram Research, Inc., “Mathematica, Version 12.1,” Champaign, IL, 2020.
- [174] J.-F. Paquet and S. A. Bass, *Phys. Rev. C* **102**, 014903 (2020), arXiv:1912.06287 [nucl-th].
- [175] https://github.com/j-f-paquet/ideal_hydro_cylindrical.
- [176] <https://github.com/Duke-QCD/frzout/tree/61b50f627174d8696f33413ab1ff873c4f0c952b>.
- [177] M. Tanabashi *et al.* (Particle Data Group), *Phys. Rev. D* **98**, 030001 (2018).