Constraining physical models at gigabar pressures

J. J. Ruby [1,2] J. R. Rygg,[1,2,3] D. A. Chin [1,2] J. A. Gaffney,[4] P. J. Adrian [5] D. Bishel,[1,2] C. J. Forrest,[2] V. Yu. Glebov,[2]
N. V. Kabadi,[5] P. M. Nilson,[2] Y. Ping,[4] C. Stoeckl,[2] and G. W. Collins[1,2,3]

[1]*Department of Physics and Astronomy, University of Rochester, Rochester, New York 14627, USA*
[2]*Laboratory for Laser Energetics, University of Rochester, Rochester, New York 14627, USA*
[3]*Department of Mechanical Engineering, University of Rochester, Rochester, New York 14627, USA*
[4]*Lawrence Livermore National Laboratory, Livermore, California 94550, USA*
[5]*Plasma Science and Fusion Center, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*

High-energy-density (HED) experiments in convergent geometry are able to test physical models at pressures beyond hundreds of millions of atmospheres. The measurements from these experiments are generally highly integrated and require unique analysis techniques to procure quantitative information. This work describes a methodology to constrain the physics in convergent HED experiments by adapting the methods common to many other fields of physics. As an example, a mechanical model of an imploding shell is constrained by data from a thin-shelled direct-drive exploding-pusher experiment on the OMEGA laser system using Bayesian inference, resulting in the reconstruction of the shell dynamics and energy transfer during the implosion. The model is tested by analyzing synthetic data from a one-dimensional hydrodynamics code and is sampled using a Markov chain Monte Carlo to generate the posterior distributions of the model parameters. The goal of this work is to demonstrate a general methodology that can be used to draw conclusions from a wide variety of HED experiments.

I. INTRODUCTION

Large-scale, high-energy-density (HED) experimental facilities are able to generate states of matter that push beyond the limits of where most physical models were developed. This necessitates quantitative experiments to guide theories and techniques at such conditions. Convergent geometries, either spherical or cylindrical, amplify the pressure generated from HED drivers, such as lasers or pulsed-power machines, and create the most extreme thermodynamic states currently achievable in the laboratory. These convergent experiments, however, complicate measurements, making direct measurement of state variables and transport properties difficult. Often measurements are the result of an integrated system, where many properties of the system are responsible for observations, rather than a small number of key physical quantities. Such integrated system analysis requires a new way of approaching how experimental data are interpreted.

The study of inference methodology is an established field [1] and modern techniques are regularly used in many areas of physics [2]. Although various methods of inference have successfully been used to understand HED experiments [3–7], these methods are not widely implemented. Use of these techniques has provided insight across many fields of physics [2,8] and there is great opportunity for discovery in HED science by adopting these methods.

This work proposes the use of a well-established [2,8] methodology for deriving quantitative physical information from integrated HED experiments. This process includes the synthesis of a model, which contains the essential physics

from the experiment, and is used in a forward analysis, generating synthetic experimental data to directly compare to the measurements. Given the proper parametrization and set of measurements the model can be constrained, giving a quantitative assessment of the input parameters. An example workflow is given, using a direct-drive exploding-pusher experiment carried out on the 60-beam OMEGA laser system, where the trajectory of an exploding shell is measured via x-ray self-emission on an x-ray framing camera. The process is applied to synthetic data generated using the one-dimensional (1D) hydrodynamics code LILAC, confirming that the parametrization accurately represents the physical quantities of interest. The process is used to infer the temporal pressure profile felt by an in-flight shell and the subsequent trajectory from experimental measurements, demonstrating how a large amount of physical insight can be gained by using integrated measurements.

II. METHODOLOGY FOR MULTIVARIATE ANALYSIS OF HED EXPERIMENTS

It is useful to picture the flow of information in a typical HED experiment as shown in Fig. 1. Shannon presented the first picture of information flow in [1] shown in Fig. 1(a) and it is adapted to relate to a physical system in Fig. 1(b). The information source in these experiments is the physical system itself and the message can be thought of as the physical properties of the system. The message is transmitted through some other physical processes, usually self-emission or interaction with some kind of probe. The message is subject to noise,
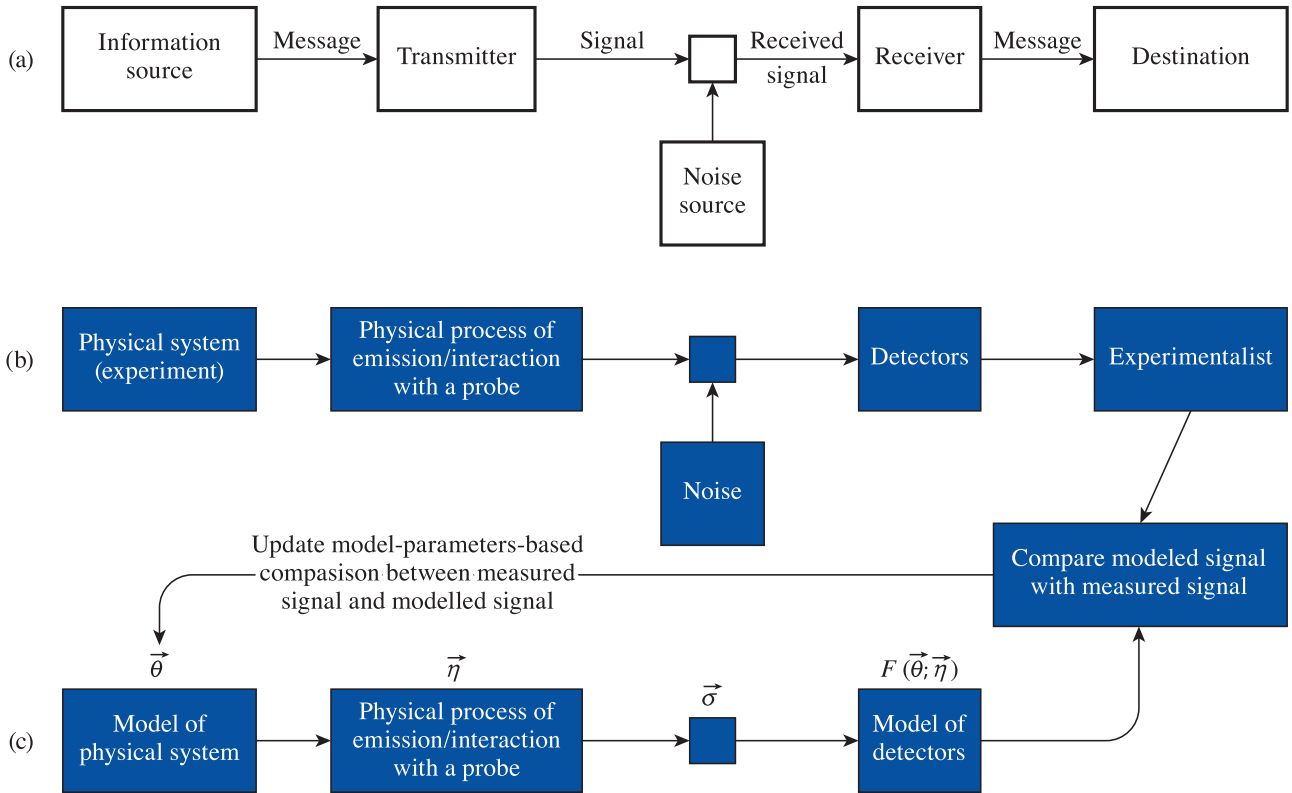
FIG. 1. (a) Flow of information in the transmission of a message as shown by Shannon in [1]. (b) Interpretation of the flow of information in a physics experiment. The information source is the physical system and the message is the details of that system. The information is transmitted either through some self-emission process or an interaction with some kind of probe. The signal is subject to noise during the transmission such as background signals that are not relevant to the details of the system that are of interest. The receiver is the detector used to measure the signal and the received message is in the form of the raw data output (e.g., an image). The destination of the signal is the experimentalist making the measurement. (c) The proposed workflow of this article in the form of a standard Bayesian forward inference. The workflow models the flow of information shown in (b) starting with a model of the physical system parametrized by the set of values $\vec{\theta}$ producing the signal through the modeled emission and interaction process parametrized by the set of values $\vec{\eta}$. The signal is then passed through a model of the detectors $F(\vec{\theta}; \vec{\eta})$, which outputs a message that is directly compared with the experimentally measured message. The model parameters are then updated based on the likelihood of the measured signal, which takes into account the experimental noise through the uncertainty $\vec{\sigma}$.

such as unwanted background processes, and is received by detectors. Finally, the destination is the experimentalist and the received message comes in the form of raw data, such as images or oscilloscope traces. Similar to how Shannon sought to take a noise-encoded signal and reconstruct the original message, the experimentalist has the same problem. The raw data can be thought of as an encoded message that is subject to noise through transmission.

In this article, a method to reconstruct the original message is presented based on Bayesian inference. A general forward-modeling procedure is shown in Fig. 1(c), where a physical model of the system, parametrized by the set of quantities $\vec{\theta}$, is used to generate a message by modeling the process of self-emission and/or probe interaction, which is parametrized by the set of quantities $\vec{\eta}$. The signal is subject to some noise, which is represented by $\vec{\sigma}$, the uncertainty on the measure signal. Finally, detectors are modeled by some function $F(\vec{\theta}; \vec{\eta})$ that creates a synthetic signal that can be directly compared to measured signal.

At this point, the comparison between the synthetic signal and measured signal can happen with any choice of metric, such as the sum-squared pointwise deviation (as in least squares analysis), and the analysis does not have to take the form of Bayesian inference. This process can be done by minimizing the chosen metric [4], for example, to get the "best fit" values for $\vec{\theta}$.

Although a Bayesian method of inference is not necessary in general, there are some cases where it is necessary, and in all cases it has significant benefits over other methods of model fitting. These advantages include a better understanding of uncertainties and constraint, an intuitive and robust methodology for model comparisons, and a straightforward method for including additional information in the model fitting. A list of common terms of Bayesian inference with definitions that will be used throughout this work is given in Table I.

In the highly integrated systems that are commonly found in HED science, Bayesian inference addresses many difficulties found in analysis. Perhaps the most important advantage comes from the full construction of the posterior probability distributions, usually through Markov chain Monte Carlo techniques, which not only give insight into the uncertainties

TABLE I. Common terms and definitions found in Bayesian inference. $p(X|Y)$ represents the probability distribution of X given that Y is true.

Term	Mathematical formulation	Description
Model	M	M_1 or M_2
Data	D or $\vec{Y}$	Observed data
Model parameter	$\vec{\theta}$	Parameters in the model
Parameter posterior distribution	$p(\vec{\theta}|D, M) = \frac{p(D|\vec{\theta}, M)p(\vec{\theta}|M)}{p(D|M)}$ $p(\theta_i|D, M) = \int d\theta_{j \neq i}\, p(\vec{\theta}|D, M)$	Probability distribution of parameters given data
Likelihood	$p(D|\theta_i, M)$	Probability of observed data given a model and set of parameters
Parameter prior distribution	$p(\theta_i|M)$	Represents initial information about the parameter
Model posterior	$p(M_1|D) = \frac{p(D|M_1)p(M_1)}{p(D)}$	Probability of model given observed data
Marginal likelihood	$p(D|M_k) = \int d\vec{\theta}\, p(D|\vec{\theta}, M_k)\, p(\vec{\theta}|M_k)$	Probability of observed data given a model integrating over parameter distributions
Bayes factor	$B = \frac{p(D|M_1)}{p(D|M_2)}$	Model comparison metric assumes $p(M_1) = p(M_2)$

on any given parameter, but also capture multiple explanations for the observed phenomena. This is critically important because in these systems there is no guarantee that the observations are sufficient to uniquely infer the parameters, and this knowledge would be lost when using a point estimate technique. The construction of the full posteriors also gives a natural way to compare different models using the marginal likelihood or other information criteria as discussed in Sec. V D. Finally, Bayesian inference includes the use of prior distributions for the values in $\vec{\theta}$ and $\vec{\eta}$, allowing previous measurements, physical constraints, and any other relevant information to be used to constrain a model that would otherwise be underdescribed by a single measurement.

It should be stressed again that the ideas presented here are not, in themselves, novel and have been well developed over a long period of time and successfully used in many fields [2,8]. The contribution of this work is to adopt the methods used elsewhere and adapt them for use in HED systems. This has been suggested before [5] in the context of inertial confinement fusion measurements but has yet to be widely used in HED analysis. These methods will provide an avenue for using integrated measurements, which previously were very challenging or impossible to analyze, to gain quantitative insight into physics in extreme conditions.

III. EXPERIMENTAL MEASUREMENTS

The experiment that will be used as an example for the application of this method was performed on the 60-beam OMEGA laser system [9] at the Laboratory for Laser Energetics. A sketch of the experimental setup is shown in Fig. 2. The target is an 879-μm SiO_2 glass sphere with 3-μm-thick walls. The sphere was filled with 18.9 atm of deuterium gas. The target was illuminated with 60 spherically configured uv laser beams containing 230 J per beam in a 600-ps super-Gaussian square pulse. The beams were configured with SG-5 distributed phase plates [10], resulting in a super-Gaussian spatial profile with diameter of about 720 μm.

The lasers drive a strong spherically converging shock wave into the glass shell via mass ablation [11]. This shock transits the glass shell and breaks out into the deuterium gas, eventually reaching the center of the target. Once at the center, the shock is strong enough to generate deuterium-deuterium fusion reactions and a nuclear particle yield is measured. The shock then moves back outward through the deuterium gas and eventually interacts with the glass shell that was set on a converging trajectory by the first shock interaction. Once the rebound shock hits the shell, the material becomes hot enough to have significant ionization and produces bremsstrahlung x-ray emission. This emission is measured using a pinhole array that projects two-dimensional (2D) images onto an x-ray framing camera that uses a voltage sweep to temporally gate the images [12]. This gives a series of snapshots of the emitting shell projected onto the 2D camera. The center of each snapshot is found, and a radially averaged measurement of the x-ray emission is taken, shown in Fig. 3(b). Since the shell emission is localized in space, this measurement gives the trajectory of the shell, which will be used to constrain the model presented in this work. The trajectory is shown in Fig. 3(c) with the corresponding errors determined from the experimental uncertainty in magnification and from the variation in the radially averaged location of the peak emission.

IV. MODEL SYNTHESIS

Analysis of physical data brings some model of the underlying system to bear, whether it be implicitly or explicitly. Implicit assertion of a model sometimes occurs when a direct correspondence is drawn between some measured quantity and some physical quantity of interest. Often, it is preferable
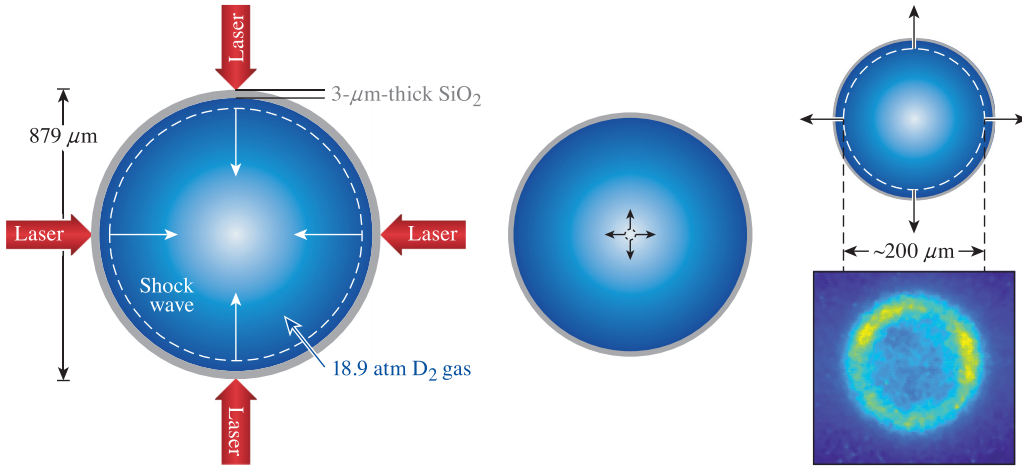
FIG. 2. A cartoon of the experimental configuration. UV laser beams uniformly irradiate an SiO₂ shell, ablate some portion of the shell mass, and drive a strong spherical shock wave. The shell is set in motion while the shock converges in the center of the target, producing nuclear fusion products. The shock then diverges from the center eventually reaching the shell, now at about $\frac{1}{4}$ the initial radius. When the shell is hit by the rebounding shock wave, the temperature is sufficiently high to produce x-ray emission via bremsstrahlung radiation, which is measured using an x-ray framing camera (example shown at bottom right).

to explicitly impose the underlying model assumptions and, when possible, forward model the corresponding measurements so as to keep the data as close to original as possible. Each measurement and system is unique and the point that the model should be compared to the data must be carefully considered; inevitably, some level of processing will occur to both the observed and modeled data. An example of this decision is whether to define error bars for the experimental data, as is done in this work, or apply noise with an appropriate distribution to the simulated data, which could be used, for example, in the processing of spectral measurements or image comparison where a pointwise error estimation may be cumbersome. Ultimately, this decision is made on a case-by-case basis and depends on the nature of both the data and the model.

A. Desired model properties

This work intends to quantitatively determine meaningful measurements from experimental data. This end requires that

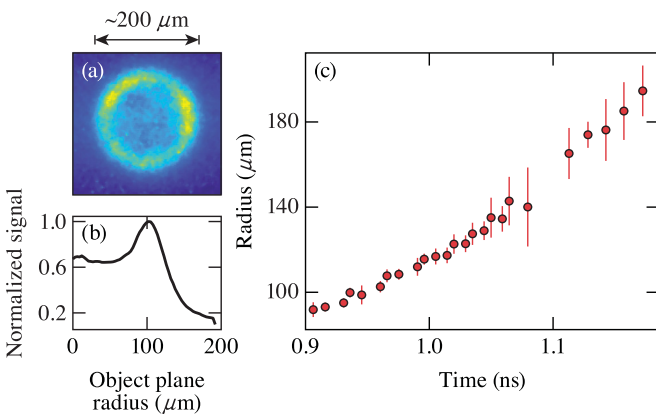FIG. 3. (a) An example frame of the emitting shell, (b) the radially averaged line out, and (c) the measured trajectory of the diverging shell.

whatever model used is parametrized by physically meaningful quantities, so when a fit of these parameters is made to experimental data the results are able to be abstracted beyond the model. There is no guarantee of uniqueness when constructing a model. When faced with multiple models that adequately explain the observations' preference is given based on Ockham's razor [8,13], that is the simpler of the models is preferred. This can be evaluated quantitatively in multiple ways [2,8], for example, by making use of Bayes factors [13,14] or other information criteria [15]. Additionally, it is desirable for the model to be computationally efficient enough to be sampled for Bayesian inversion, especially if there are multiple candidate models that need to be compared. This is not always possible, and depends on the specific requirements and complexities of the system; even when not possible directly there is potential for a surrogate model [16,17] to be used in order to regain the computational efficiency needed to invert the model with a Markov chain Monte Carlo (MCMC). The model presented in Sec. IV B, which required on the order of 1×10^7 evaluations per sampling, does not require a surrogate. The sampling of the inversion is discussed in Sec. V C.

B. Model for exploding pusher targets on OMEGA

The model used to describe the experiment discussed in Sec. III uses 10 total parameters. There are three parameters that describe the shell, M_f, R_f, and v_f, which are the mass of the shell, radius of the shell, and velocity of the shell, respectively. The second-order ordinary differential equation (ODE)

$$\frac{d^2 R}{dt^2} = a, \tag{1}$$

where R is the radial location of the shell and a is the acceleration of the shell. The acceleration takes the form

$$a(t) = \begin{cases} 0, & t < t_a \\ \frac{4\pi R^2 P}{M} - \frac{\dot{M} v_e}{M}, & t_a \leqslant t < t_l \\ \frac{4\pi R^2 P}{M}, & t \geqslant t_l \end{cases} \tag{2}$$

where P is the pressure pushing outward on the shell due to the fuel inside and M is the mass of the shell both of which (along with the radius R) are time-dependent quantities, while the ablation rate $\dot{M}$ and exhaust velocity v_e are both time independent. The acceleration starts when $t \geqslant t_a$, interpreted as the time ablation starts. At this point, there is an outward force from the fuel inside the shell and an inward force from the rocket effect of the ablated mass. Once the laser turns off at t_l the force from mass ablation stops and only the outward force from the fuel remains. The pressure profile from the fuel is given by

$$P(t) = \begin{cases} (P_{rs})e^{-\gamma_g(t_s-t)}, & t \leqslant t_s \\ (P_{rs})e^{-\gamma_d(t-t_s)}, & t \geqslant t_s \end{cases} \tag{3}$$

with exponentially growing term with growth rate γ_g, for times before t_s and an independent exponentially decaying profile with decay rate γ_d for times after t_s. The peak pressure reached is given by P_{rs}, which is interpreted at the pressure of the rebounding shock wave.

The mass of the shell is given by

$$M(t) = \begin{cases} M_0 - \dot{M}(t - t_a), & t \leqslant \min[t_s, t_M] \\ M_{min} + \dot{M}(t - \min[t_s, t_M]), & t > \min[t_s, t_M] \end{cases} \tag{4}$$

where t_M is the time it takes to lose all of the mass of the shell and M_{min} is the minimum mass the shell achieves. The shell continually loses mass to ablation at a constant rate while the laser is on and to material release after the laser stops until the shell trajectory reaches its minimum. At this point, the shell starts moving outward and regains the released mass as it moves out.

V. MODEL INFERENCE

The data discussed in Sec. III are used to constrain the model. This data set consists of $N = 26$ data points in the form of radius and time pairs that correspond to the peak of the radially averaged emission measured on the framing camera. These data points, along with some additional knowledge such as target metrology, measurements of the laser pulse, and physical limits, were used to constrain the model.

A. Objective function

In order to perform parameter estimation, an objective function must be used. In this work the log-likelihood function $\ln(\mathscr{L})$,

$$\ln(\mathscr{L}) = -\frac{1}{2} \sum_{i=0}^{N} \left[\frac{Y_i - F(\vec{\theta})_i}{\sigma_i} \right]^2, \tag{5}$$

is used where the Y_i are the set of experimental measurements with uncertainties given by σ_i, the $F(\vec{\theta})_i$ are the set of modeled data values that follow from the input parameter values $\vec{\theta}$ being used for the model

$$\vec{\theta} = [M_0, R_0, \dot{M}, v_e, t_a, t_l, t_s, P_{rs}, \gamma_g, \gamma_d], \tag{6}$$

so the function $F(\vec{\theta})$ takes in a set of values for each of the parameters, constructs the pressure profile and mass profile given by Eqs. (3) and (4), respectively, and solves the ODE

given by Eq. (1). The function then returns the radius of the shell at the experimentally measured times to be directly compared to the measured values.

This formulation assumes independent Gaussian distributed errors on the experimental measurements, but that is not necessary and the distribution of the measurements can take any form and a corresponding likelihood function can be defined.

Parameter estimation can be done using the log-likelihood in the form of a maximum likelihood estimate (MLE) or, equivalently, a minimum χ^2 analysis, but these methods are point estimates and have no notion of uncertainty. This work uses Bayesian inference, constructing the full posterior distribution of the parameters using an MCMC sampling. Having access to the posterior distributions of the parameters provides much greater insight into the behavior of the model and the constraint new data has on the parameters along with the uncertainty and correlations among these parameters that point estimates, such as MLE, do not provide. Uncertainty analysis using a minimum χ^2 analysis usually consists of finding a minimum (which has no guarantee of being the global minimum in nontrivial cases) and then seeking out the $\chi^2_{min} + 1$ contours and using those as confidence intervals on parameters. This process implicitly makes assertions about the shape of the likelihood surface and treats the likelihood as a probability density in the parameters, which is mathematically inconsistent [18].

In Bayesian inference the likelihood is multiplied by the prior probability on the parameters and integrated, resulting in a proper probability density that can be integrated to produce confidence intervals. This process is not subject to any individual point on the likelihood surface and is therefore capable of dealing with complicated and multimodal surfaces that would hinder MLE techniques.

B. Model testing with synthetic data

Whether a model has a physically relevant parametrization and whether or not it is constrained given a particular measurement must be addressed before using a model to gain insight into the physics of a system. This is done by using a simulation of an experiment, in this case provided by the 1D Lagrangian hydrodynamics code LILAC [19], and producing synthetic experimental data, which is then fit by the model. The underlying physics in the simulation can then be scrutinized to see how well the model parametrization compares. The ability to simulate both physical systems of interest and detectors used to measure those systems is a critical aspect of Bayesian inference and provides great insight into not only analysis, but also the design of experiments.

The LILAC simulation was run using experimental spatial and temporal laser beam profiles, a flux-limited thermal-transport model with flux limiter $=0.12$, a $3-\mu$m SiO_2 shell using SESAME equation of state 7380 [20] (for reference more recent tables include SESAME 7360 and 7361), and 436.5 μm of deuterium gas at initial density $= 10.4$ mg/cm^3 (pressure $= 18.9$ atm) using SESAME equation of state 5262 [20]. The simulation outputs were postprocessed to calculate the x-ray emission in each zone by calculating

the bremsstrahlung emissivity [11,21,22] given by

$$J_\nu = \frac{\sqrt{2^{11}\pi^3}}{3} \frac{\left(\frac{e^2}{4\pi\epsilon_0}\right)^3 \bar{Z}^2 n_e^2}{m_e^2 c^3 \bar{Z}\sqrt{3kT_e/m_e}} e^{-\frac{h\nu}{kT_e}} \quad [Js^{-1}\,\mathrm{m}^{-3}\,\mathrm{Hz}^{-1}],$$

(7)

where J_ν is the spectrally dependent emissivity; $\bar{Z}$ is the average ionization state (here $\bar{Z} = 1$ for deuterium and $\bar{Z} = 10$ for SiO$_2$); n_e and T_e are the electron density and temperature, respectively; ν is the frequency of the emitted photon; and all other symbols are the usual physical or mathematical constants. The emissivity of each zone is then multiplied by the zone volume, integrated over 30 ps, and integrated for photon energies above 5 keV. The temporal and spectral integrations both are meant to roughly represent the temporal and spectral response of the framing camera used in the experiment. The result of this is a series of radial line outs of x-ray emission from the simulation, which are then Abel transformed in order to replicate the signal of the projected 2D image measured in the experiment.

C. Posterior sampling of the model

The posterior distributions were sampled using a sequential Monte Carlo (SMC) [23], from the PYMC3 [24] library, which samples a tempered distribution using, for example, a Metropolis Hastings algorithm. As the "temperature" increases, the tempered distribution converges to the posterior distributions. SMC methods have the advantage of much more efficient sampling of higher-dimensional spaces, in this case a 10-dimensional space, and they are also better at sampling from potentially multimodal distributions. An additional benefit of SMC sampling is that the prior likelihood surface is fully sampled as the first step in the SMC, providing a necessary piece for calculating marginal likelihoods, a valuable metric for model comparisons.

A sampling for the simulated data $\vec{Y}_{\mathrm{sim}}$ was done using the likelihood function in Eq. (5) using 3-μm error on each of the points. Prior selection is an important aspect of Bayesian inference for parameter estimation and especially for model comparison (see Sec. V D). The choice of priors is one way that additional information, outside of measured data, can be used to constrain an otherwise underdescribed model. In this case, there are 10 parameters in the model in total, but 5 of the parameters have strong priors due to physical limitations or measurement. These five parameters are R_0, M_0, t_a, t_l, and t_s. The first two are target parameters that are metrologized prior to the experiment so they take the form of normal distributions around the central values with uncertainty given metrology limitations. The last three parameters all coincide with temporal events observed in the experiment, such as the laser turning on and off, and the onset of shell emission corresponding to shock heating. Additionally, physical considerations such as the mass, radius, and pressure all being strictly positive provide constraint. The other parameters, which are not as well constrained *a priori*, receive broad normal priors that are mostly uninformed with the exception of physical of bounds such as for the ablation rate and exhaust velocity both of which must be strictly positive quantities. The Bayesian framework allows all of this information to be explicitly imposed and uses them to further constrain the system.
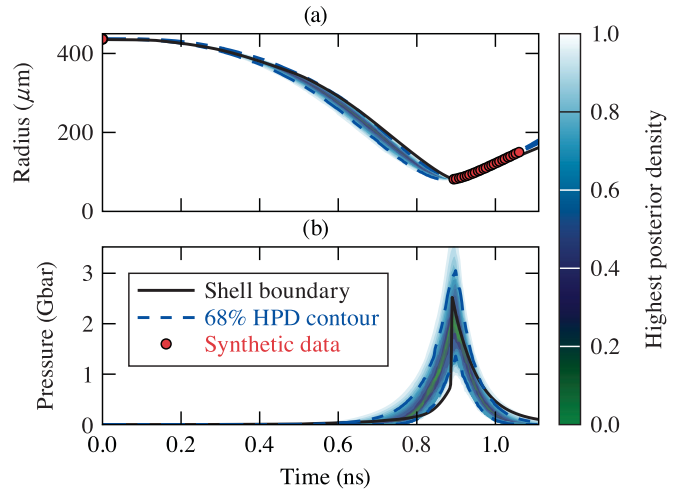

FIG. 4. The posterior predictive distributions for (a) the shell trajectory and (b) the temporal pressure profile at the shell fuel interface based on sampling the posterior distributions of the model parameters estimated using the synthetic data generated by LILAC. Also shown is the (a) trajectory of the innermost SiO$_2$ zone in LILAC and (b) the pressure in that zone, both shown in black, and the synthetic data points in red (a). The color scale shows highest posterior density intervals (HPD), which give the smallest region that contains a given amount of probability (dotted-dashed blue lines represent the 68.3% probability HPD). Although the model is only constrained by the initial radius and the data points starting at near 0.9 ns, it is able to predict the position of the inner edge of the shell and pressure the fuel exerts on the shell at all times in-between. The greatest divergence happens in the late time directly before the observations where both the trajectory and pressure profile fall outside the 68.3% prediction band.

Once the posterior distributions of the model parameters are constructed (through the SMC sampling) values can be drawn from the distributions and passed through the model to generate posterior predictive distributions (PPD) for the shell trajectory and pressure profile. The PPD gives a prediction for what additional measurements would provide, based on the given model. In this case, for a given time there is a probability distribution for the radius and the pressure and by sampling the PDD at many times a probabilistic band of trajectories and pressures profiles can be constructed (see Fig. 4). This process is the same as described in Sec. V A but rather than the parameters in $\vec{\theta}$ taking a single value, they each have a distribution, and instead of $F(\vec{\theta})$ returning the trajectory at the experimentally measured points the trajectory is returned at an arbitrary time of interest.

Since the synthetic data were generated from a hydrodynamics code, the trajectory of the shell and pressure profile is known for all times. In this case, comparison is made to the innermost zone of SiO$_2$ in the code, which is the zone that borders the fuel. The trajectory and pressure in this zone are shown as solid black lines in Fig. 4. Intervals of increasing probability for the PPD are shown by the color scale. The model displays good predictive capability across all times, with the "true" values from the code usually falling inside the 68.3% confidence interval (dotted-dashed blue lines). This demonstrates that the model has predictive capability outside of the regions where data are directly measured.
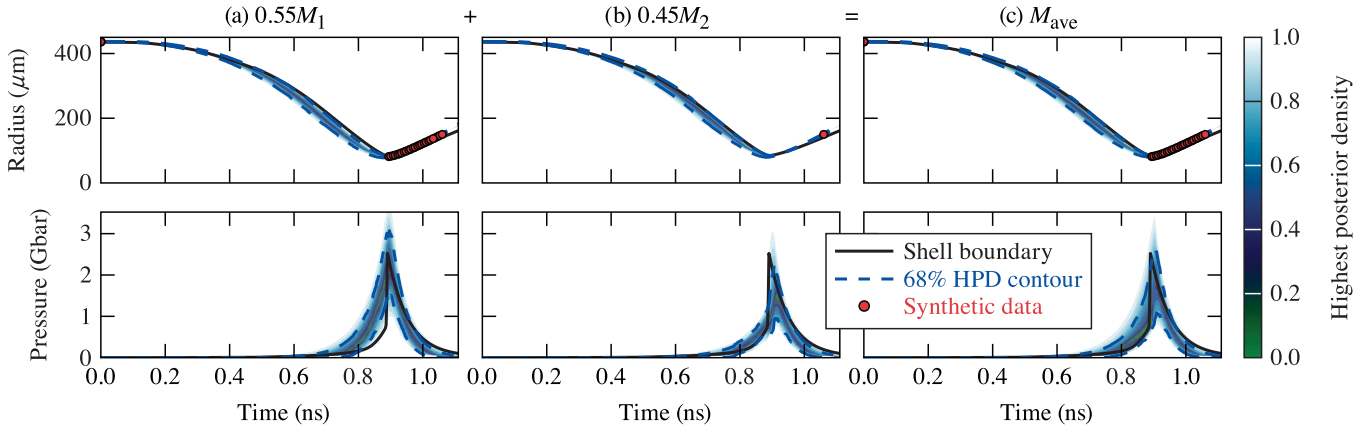
FIG. 5. PPD for the trajectory and pressure profiles for (a) M_1, (b) M_2, and (c) the averaged model based on weightings from the WAIC. The weights for M_1 and M_2 are given above the images and are 55% and 45%, respectively. The averaged model is constructed by taking the samples from each model based on the relative weightings. The combined model shows improved agreement with the simulation trajectory falling within the 68.3% HPD interval at all times and better agreement in both the preshock pressure profile than M_1 and around the time of peak pressure than M_2.

D. Comparison of different models

As discussed in Sec. IV A, there is not a unique model to describe the system. It is important to consider a range of models and quantify how they compare. Differences in model can come in the form of different functional forms $F(\vec{\theta})$ or simply a different choice of priors on models that otherwise have the same expressions. Here, an example of a comparison between two different models with different functional forms is given. M_1 is the model from Sec. IV B, and model M_2 has identical parameters except for the pressure profile, which has a different amplitude before and after t_s:

$$P(t) = \begin{cases} (P_0)e^{-\gamma_g(t_s-t)}, & t \leqslant t_s \\ (P_1)e^{-\gamma_d(t-t_s)}, & t \geqslant t_s. \end{cases} \quad (8)$$

This profile is meant to represent the jump in pressure due to the shock front hitting the shell. This model is interesting because it may be able to account for the region where the PPD of the pressure performs the worst, as seen in Fig. 4. Although it may seem obvious that a model with a pressure jump would be more appropriate, there is no guarantee that the measurements, which happen in trajectory space and after the shock hits the shell, have leverage on the pressure profile earlier in time. This coupled with the additional parameter for pressure means the M_2, while physically motivated, may not be preferred over M_1.

Model comparison is standard practice in Bayesian inference, and there are many methods that can be used. In this case, three different methods were compared: widely applicable information criteria (WAIC) [25,26], leave one out cross validation (LOO) [25,26], and Bayes factor [15,27] calculations.

Unlike model comparisons based on point estimates of parameters, such as likelihood ratios or F-test metrics, Bayesian model comparisons account for the full posterior distribution of parameters. This means that the uncertainty in parameter estimates is also part of the calculation [26]. The Bayesian information criteria (WAIC and LOO) attempt to quantify the predictive capability of the model to unobserved data, while comparing Bayes factors amounts to a more-generalized form

of a likelihood ratio test that includes marginalizing over the posterior distributions of each parameter, therefore not relying on a point estimate.

All three metrics agree when comparing M_1 and M_2, showing a slight favor for M_1. A benefit of Bayesian model comparison is that it provides a weighting to average the models together.

In this case, PPD samples of the trajectory and pressure profiles are sampled based on the weights assigned by the information criteria, resulting in PPD's for the trajectory and pressure profile based on both models. An example composite model is shown in Fig. 5, where the trajectory and pressure profile shown in Fig. 5(c) are composed of a weighted sampling of the model, M_1 and M_2, shown in Figs. 5(a) and 5(b), respectively. In this case, the weighting was established by using the WAIC information criteria, but LOO and the Bayes factors gave very similar weightings.
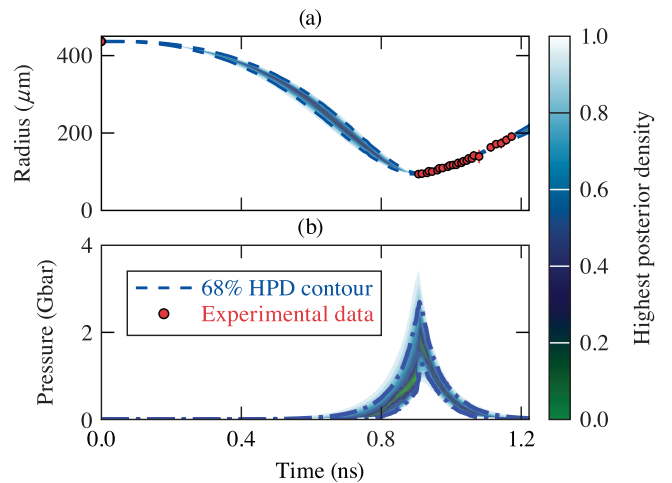

FIG. 6. (a) The shell trajectory and (b) pressure profile for the model sampled using the experimental data. The model values are shown as regions of confidence given by the HPD. The 68% HPD contour is given by the blue dashed curve.
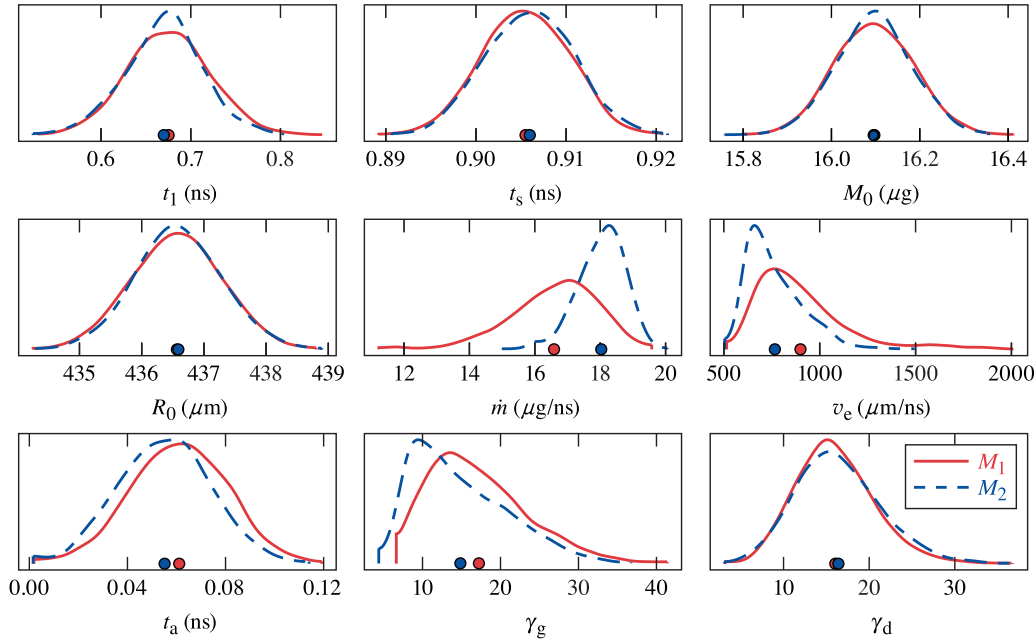
FIG. 7. The posterior probability density distributions resulting from the MCMC sampling for the nine parameters that are common to both model M_1 (red) and M_2 (blue). The mean values of each distribution are given by the points on the bottom axis. The only parameters with a significant deviation are $\dot{M}$ and v_e, due to how they are related to the pressure profile as discussed in Sec. VI and shown in Fig. 9.

VI. RESULTS

M_1 and M_2 were sampled as in Sec. V, but now using the experimental data $\vec{y}_{\text{expt}}$. The models were then averaged the same way, with relative weightings of 0.54 and 0.46, respectively (given by WAIC), and the resulting trajectory and pressure profile is shown Fig. 6.

The posterior distributions of the parameters give important insight into how well constrained the models are by the data. Figure 7 shows the posterior distributions for the nine parameters that are common to both M_1 and M_2. Among these, most have similar distributions except for $\dot{M}$ and v_e. The difference in these distributions, specifically the fact the they are more peaked for the M_2 case, is due to how the pressure profile is treated early in time. The posteriors for the peak pressures, which are not common between the two models, are shown in Fig. 8.

Figure 9, displaying pairwise correlations between parameters, shows a strong correlation between $\dot{M}$, v_e, and γ_g, which is physically intuitive because the trajectory of the shell is set by a balance between the outward pressure from the gas (determined by γ_g) and the ablation pressure (determined by a combination of $\dot{M}$ and v_e). It was shown in Fig. 4 that M_2 is better at constraining the early-time pressure profiles and therefore would have less uncertainty in the correlated parameters $\dot{M}$ and v_e, leading to more peaked distributions.

Another benefit of constructing the full posteriors and examining pairwise correlations is the opportunity to identify methods of constraining otherwise difficult to measure variables [28]. Since $\dot{M}$ and v_e are tightly correlated, the measurement of either the mass ablation rate or the exhaust velocity of the ablated material would tightly constrain the other value. Likewise, a measurement of either of these would also be very constraining of the pressure exerted on the shell by the fuel or, conversely, a measurement of the pres-

sure exerted on the shell would be very constraining on the mass ablation rate and exhaust velocity. Examining one self-consistent model of the entire experimental system provides insights that would otherwise be difficult to identify (although in retrospect seem obvious).

The methods presented here for analyzing these types of experiments provide an avenue for recovering as much information as possible from the data that are measured. There is a huge depth of knowledge waiting in the message (referencing back to Sec. I) and this procedure provides an avenue to reconstruct this message with as much fidelity as possible.
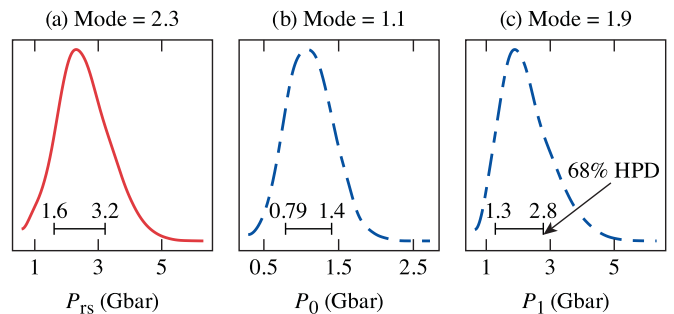

FIG. 8. The posterior probability density distributions resulting from the MCMC sampling for the three different pressure parameters P_{rs} of M_1 (red curve) and P_0 and P_1 of M_2 (blue curve). Also shown are the 68% HPD intervals which denote the shortest interval that contains 68.3% of the probability in the distribution and the values of the mode (peak) of the distributions. The values of P_{rs} and P_1 can both be interpreted as the pressure of the shock wave when it interacts with the shell material and P_0 can be interpreted as the pressure just before the shock reaches the shell in the two-pressure model (M_2).
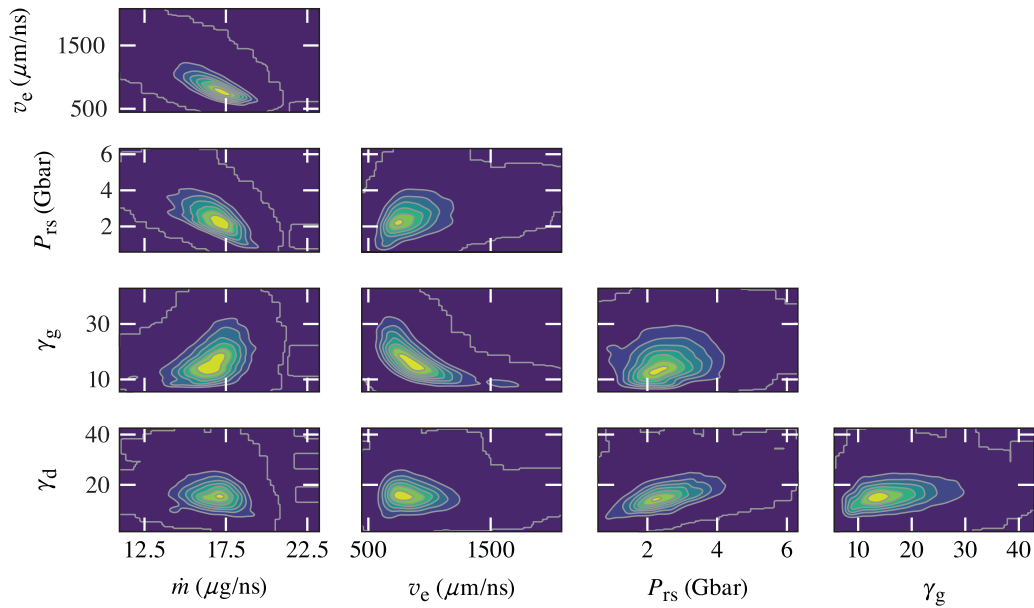
FIG. 9. Kernel density estimate of the pairwise posterior distributions of parameters without strong prior information (as discussed in Sec. V C). Only the pressure parameter from M_1 is shown, P_{rs}, as the information is redundant for the other pressure parameters. Pairwise correlations give insight into physical connections between variables and show how one variable can be constrained through measurement of another. An example is the negative correlation between $\dot{M}$ and v_e (discussed in Sec. VI), meaning that a constraint of the exhaust velocity would greatly constrain the mass ablation rate in a laser-driven shell experiment. These types of connections can offer insight into how to constrain otherwise difficult to measure quantities.

Additional insight into the physical results of this process is present in Paper I [29].

VII. CONCLUSION

The methods presented showed how a seemingly isolated, integrated measurement of an evolving system can be used to reconstruct a model of the entire system with enough fidelity to gain quantifiable physical insight. This methodology has a strong history of being used in other fields of physics and has the potential to greatly impact HED physics by making possible quantified measurements in conditions otherwise inaccessible. The combination of new analysis techniques for integrated measurements with high-quality, existing experimental methods and facilities capable of creating conditions otherwise impossible on Earth promise to provide new insight into our understanding of physics at these extreme conditions.

ACKNOWLEDGMENTS

[1] C. E. Shannon, Bell Syst. Tech. J. **27**, 379 (1948); **27**, 623 (1948).

[2] U. von Toussaint, Rev. Mod. Phys. **83**, 943 (2011).

[3] D. Hicks, B. Spears, D. Braun, R. Olson, C. Sorce, P. Celliers, G. Collins, and O. Landen1, Phys. Plasmas **17**, 102703 (2010).

[4] C. Cerjan, P. Springer, and S. Sepke, Phys. Plasmas **20**, 056319 (2013).

[5] J. Gaffney, D. Clark, V. Sonnad, and S. Libby, Nucl. Fusion **53**, 073032 (2013).

[6] P. Bremer, D. Malijovec, A. Saha, B. Wang, J. Gaffney, B. Spears, and V. Pascucci, Comput. Visualization Sci. **17**, 1 (2015).

[7] M. F. Kasim, T. P. Galligan, J. Topp-Mugglestone, G. Gregori, and S. M. Vinko, Phys. Plasmas **26**, 112706 (2019).

[8] V. Dose, Rep. Prog. Phys. **66**, 1421 (2003).

[9] T. Boehly, D. Brown, R. Craxton, R. Keck, J. Knauer, J. Kelly, T. Kessler, S. Kumpan, S. Loucks, S. Letzring *et al.*, Opt. Commun. **133**, 495 (1997).

[10] T. Kessler, Y. Lin, J. Armstrong, and B. Velazquez, Proc. SPIE **1870**, 95 (1993).

[11] S. Atzeni and J. M. ter vehn, *The Physics of Inertial Fusion Beam Plasma Interaction, Hydrodynamics, Hot Dense Matter*, 1st ed. (Oxford University Press, Oxford, 2004).

[12] J. Kilkenny, Laser Part. Beams **9**, 49 (1991).

[13] R. Rosenkrantz and E. Jaynes, J. Am. Stat. Assoc. **74**, 740 (1979).

[14] R. Kass and A. Raftery, J. Am. Stat. Assoc. **430**, 773 (1995).

[15] D. J. Spiegelhalter, N. G. Best, B. P. Carlin, and A. Van Der Linde, J. R. Stat. Soc. B **64**, 583 (2002).

[16] K. Humbird, J. Peterson, B. Spears, and R. McClarren, IEEE Trans. Plasma Sci. **48**, 61 (2020).

[17] B. Kustowski, J. A. Gaffney, B. K. Spears, G. J. Anderson, J. J. Thiagarajan, and R. Anirudh, IEEE Trans. Plasma Sci. **48**, 46 (2020).

[18] R. Cousins, Am. J. Phys. **63**, 398 (1995).

[19] J. Delettrez, R. Epstein, M. C. Richardson, P. A. Jaanimagi, and B. L. Henke, Phys. Rev. A **36**, 3926 (1987).

[20] B. I. Bennett, J. D. Johnson, G. I. Kerley, and G. T. Rood, SESAME: The Los Alamos National Laboratory Equation of State Database, Report No. LA-UR-92-3407, Los Alamos, NM, 1978 (unpublished).

[21] Y. Zeldovich and Y. Raizer, *Physics of Shock Waves and High-Temperature Hydrodynamic Phenomena* (Dover, Mineola, NY, 2002), Vol. I, Chap. 12, Sec. 2, p. 794.

[22] J. Colvin and J. Larsen, *Extreme Physics: Properties and Behavior of Matter at Extreme Conditions* (Cambridge University Press, Cambridge, 2014).

[23] A. Doucet, N. de Freitas, and N. Gordon, *Sequential Monte Carlo Methods in Practice*, 1st ed. (Springer, New York, 2001), p. 79.

[24] J. Salvatier, T. V. Wiecki, and C. Fonnesbeck, PeerJ Comput. Sci. **4**, e55 (2016).

[25] A. Vehtari, A. Gelman, and J. Gabry, Stat. Comput. **27**, 1413 (2017).

[26] A. Gelman, J. Hwang, and A. Vehtari, Stat. Comput. **24**, 997 (2014).

[27] H. Jeffreys, *The Theory of Probability* (Clarendon, Oxford, 1961).

[28] J. Gaffney, S. Brandon, K. Humbird, M. Kruse, R. Nora, J. Peterson, and B. Spears, Phys. Plasmas **26**, 082704 (2019).

[29] J. J. Ruby, J. R. Rygg, D. A. Chin, J. A. Gaffney, P. J. Adrian, C. J. Forrest, V. Yu. Glebov, N. V. Kabadi, P. M. Nilson, Y. Ping, C. Stoeckl, and G.W. Collins, Phys. Rev. Lett. **125**, 215001 (2020).