

Hierarchical Bayesian Analysis and the Preston-Tonks-Wallace Model

Michael Fugate, Brian Williams, David Higdon, Kenneth M. Hanson
James Gattiker, Shuh-Rong Chen, Cetin Unal

Abstract

We utilize data from Hopkinson-bar experiments and quasi-static compression experiments to characterize uncertainties for parameters governing the Preston-Tonks-Wallace (PTW) [1] plastic deformation model for a variety of materials. This particular plastic deformation model is designed to be valid over a range of input conditions, which include strain, strain rate and temperature. However, because of variations between experimental samples, measurement variation, as well as slight inadequacies in the model, no single parameter setting gives a good match to all of the experimental data for a given material. These deficiencies need to be taken into account when assessing the uncertainties in the model parameters. In this paper, we use a Bayesian hierarchical model to account for the variations in the experimental data. This modeling approach results in parameter estimates for each material, along with uncertainty estimates, which are the main focus of this paper. The results are summarized in the table below.

Parameter	Al	Be	DU	Ta	U-6Nb
θ	0.0364 (0.011)	0.0268 (0.00756)	0.0465 (0.0219)	0.0124 (0.00558)	0.0974 (0.0497)
κ	0.415 (0.137)	0.197 (0.112)	0.136 (0.064)	0.711 (0.156)	0.503 (0.113)
$-\log(\gamma)$	10.3 (4.0)	11.2 (5.74)	14.7 (3.03)	11.5 (1.77)	6.32 (4.16)
y_0	0.0113 (0.00323)	0.00165 (0.000248)	0.00534 (0.000766)	0.00939 (0.00129)	0.00665 (0.00351)
y_∞	0.0058 (0.000239)	0.000938 (0.000283)	0.00236 (0.000934)	0.00136 (0.00028)	0.000724 (0.000522)
s_0	0.0295 (0.0118)	0.0169 (0.00978)	0.0123 (0.00327)	0.0168 (0.00554)	0.0707 (0.0228)
s_∞	0.00781 (0.000411)	0.00337 (0.00148)	0.00485 (0.00165)	0.00295 (0.00062)	0.00646 (0.000495)

Keywords: plastic deformation model, Preston-Tonks-Wallace model, uncertainty analysis, Bayesian analysis, hierarchical model, model uncertainty, Hopkinson-bar experiments, quasi-static-compression experiments

1 Introduction

The primary sources of data that are typically used to characterize the plastic behavior of a metal are obtained in quasi-static and Hopkinson-bar experiments. In quasi-static tests, a small cylinder of the material is typically squeezed at a constant, relatively slow rate and the change in its height is measured as a function of the load on the cylinder. These measurements are easily converted to stress and strain values. In Hopkinson-bar experiments, an elastic wave is transmitted through a thin cylinder of the material and its change in dimensions measured. Although these measurements require the use of a simulation code for precise interpretation, they are straightforwardly converted to a stress-strain curve at nearly constant strain rate. The strain rates attained in Hopkinson-bar experiments are around 10^3 per second, whereas in quasi-static tests they are typically about one per second or less. We use data from these experiments to estimate uncertainties for PTW model parameters that are appropriate for dynamic simulations. The analysis of data collected from these basic experiments is a fairly straightforward nonlinear data-fitting problem.

In this report we describe results from fitting the Preston-Tonks-Wallace (PTW) model [1] to five materials: aluminum (Al), beryllium (Be), depleted uranium (DU), tantalum (Ta), and uranium 6 wt% niobium (U-6Nb). While the PTW model has been fit to some of these materials before [2], our contribution is to impose a Bayesian hierarchical model on the data and to estimate the uncertainties in parameter estimates from the Bayesian analysis. The hierarchical model is fit via Markov chain Monte Carlo (MCMC).

In Section 2 we describe the PTW model in some detail. Section 3 discusses the rationale for the hierarchical model. Section 4 describes the hierarchical model as it applies to our data. Section 5 describes the analysis of the data and results for each material type. Section 6 concludes with a discussion. The Appendix provides a more detailed presentation of the MCMC algorithm used to fit the hierarchical model.

2 Material Characterization Model

The PTW model describes the plastic deformation of metals in terms of the dependence of plastic stress on plastic strain over a wide range of strain rates and temperatures. The following summary of the PTW model is taken directly from [3].

In the PTW model, the plastic stress in a material is a function of the amount of strain ψ it has undergone, the strain rate $\dot{\psi}$, the material temperature T , and its density ρ . It is assumed that the plastic stress is independent of the history of the material. Furthermore, PTW ignores non-isotropic plasticity and material texture effects. Material fracture or failure is not incorporated in PTW.

The PTW model is written in terms of three scaled dimensionless variables. The scaled stress variable is $\hat{\tau} = \tau/G(\rho, T)$, where τ is the flow stress, which is one-half the usual von Mises equivalent deviatoric stress σ (i.e., $\tau = \sigma/2$) and $G(\rho, T)$ is the shear modulus, which is a function of the material density ρ and temperature T . The shear modulus is taken to be $G(\rho, T) = G_0(\rho) (1 - \alpha \hat{T})$, where $G_0(\rho)$ is the shear modulus at $T = 0$ and $\alpha > 0$ is a material parameter. The material temperature is scaled to its melt temperature T_m , which is a function of the material density ρ , $\hat{T} = T/T_m(\rho)$. For plastic flow, clearly $\hat{T} < 1$. The

equivalent plastic strain is denoted by ψ . The strain rate $\dot{\psi}$ is scaled to an appropriate rate

$$\dot{\xi}(\rho, T) = \frac{1}{2} \left(\frac{4\pi\rho}{3M} \right)^{\frac{1}{3}} \left(\frac{G}{\rho} \right)^{\frac{1}{2}},$$

where M is the atomic mass of the material and $\dot{\xi}$ is the reciprocal of the time for a transverse sound wave to cross an atom (i.e., atomic vibration frequency). The strain rate always appears in the PTW formulas in terms of the ratio $\dot{\psi}/\dot{\xi}$.

For any fixed values of strain rate and temperature, the scaled stress $\hat{\tau}$ ranges between the lower and upper limits given by the yield stress $\hat{\tau}_y$ and the saturation value $\hat{\tau}_s$. The functional form for $\hat{\tau}$ depends on the strain ψ as follows

$$\hat{\tau} = \hat{\tau}_s + \frac{1}{p} (s_0 - \hat{\tau}_y) \ln \left\{ 1 - [1 - \exp(-pr)] \exp \left[-\frac{p\theta\psi}{(s_0 - \hat{\tau}_y) [\exp(pr) - 1]} \right] \right\},$$

where p and θ are material-specific parameters that appear in the modified work hardening law (see [1]) and $r = (\hat{\tau}_s - \hat{\tau}_y)/(s_0 - \hat{\tau}_y)$. The parameter s_0 is explained below.

At low strain rates as a function of temperature, the plastic deformation process is controlled by thermal activation. The values for $\hat{\tau}_y$ and $\hat{\tau}_s$ are given by

$$\hat{\tau}_y = y_0 - (y_0 - y_\infty) \operatorname{erf} \left[\kappa \hat{T} \ln \left(\frac{\gamma \dot{\xi}}{\dot{\psi}} \right) \right],$$

$$\hat{\tau}_s = s_0 - (s_0 - s_\infty) \operatorname{erf} \left[\kappa \hat{T} \ln \left(\frac{\gamma \dot{\xi}}{\dot{\psi}} \right) \right],$$

where κ and γ are dimensionless material-related parameters. The error function, defined as $\operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x \exp(-t^2) dt$, has the limiting values $\operatorname{erf}(0) = 0$ and $\operatorname{erf}(\infty) = 1$. The parameters y_0 and y_∞ are the values that $\hat{\tau}_y$ takes at zero temperature and very high temperatures, respectively; s_0 and s_∞ have analogous interpretations for $\hat{\tau}_s$. Table 1 gives constraints on the parameters of the PTW model.

Parameter Constraints
$\theta > 0$
$\kappa > 0$
$\gamma \geq \dot{\psi}/\dot{\xi}$
$y_0 > y_\infty > 0$
$s_0 > s_\infty > 0$
$s_0 \geq y_0$
$s_\infty \geq y_\infty$

Table 1: PTW parameter constraints.

The PTW model is designed to extend the range of normal plastic-flow models to very high strain rates, above 10^8 s^{-1} , in which regime it relies on Wallace's theory of over-driven shocks in metals [4]. This extension introduces three additional material-specific parameters: β , y_1 , and y_2 . Because quasi-static and Hopkinson-bar experiments, which are the focus of

the present study, do not reach these very high strain rates, the formulas that apply in that regime are not given. Suffice it to say that the PTW parameters β , y_1 , and y_2 have no effect in the lower strain rate regime.

Table 2 gives the mean and standard deviation of the PTW parameters from our hierarchical analysis for each material. The posterior distribution that these summaries are based on accounts for the entire range of uncertainty in the PTW parameters across the spectrum of data sets analyzed. It is given explicitly in equation (6). The entries in the table were computed from the posterior samples by using every other realization from 100,001 to 600,000.

Parameter	Al	Be	DU	Ta	U-6Nb
θ	0.0364 (0.011)	0.0268 (0.00756)	0.0465 (0.0219)	0.0124 (0.00558)	0.0974 (0.0497)
κ	0.415 (0.137)	0.197 (0.112)	0.136 (0.064)	0.711 (0.156)	0.503 (0.113)
$-\log(\gamma)$	10.3 (4.0)	11.2 (5.74)	14.7 (3.03)	11.5 (1.77)	6.32 (4.16)
y_0	0.0113 (0.00323)	0.00165 (0.000248)	0.00534 (0.000766)	0.00939 (0.00129)	0.00665 (0.00351)
y_∞	0.0058 (0.000239)	0.000938 (0.000283)	0.00236 (0.000934)	0.00136 (0.00028)	0.000724 (0.000522)
s_0	0.0295 (0.0118)	0.0169 (0.00978)	0.0123 (0.00327)	0.0168 (0.00554)	0.0707 (0.0228)
s_∞	0.00781 (0.000411)	0.00337 (0.00148)	0.00485 (0.00165)	0.00295 (0.00062)	0.00646 (0.000495)

Table 2: Posterior mean (standard deviation) from hierarchical model analysis.

3 Why Hierarchical Modeling?

The goal of this effort is not to come up with improved “best estimates” of the PTW parameters for various materials, but rather to use the available experimental data to specify uncertainties for the parameters that are appropriate for dynamic simulations. For example, when simulating a Taylor cylinder test, the material goes through a range of temperatures and strain rates that may be orders of magnitude different from those in Hopkinson-bar experiments. We therefore would like to determine ranges for the PTW parameters that are consistent with the experiments and are wide enough to cover conceivable settings that are appropriate when temperature, strain, and strain rate may vary substantially over the course of the experiment.

In the left-hand plot of Figure 1 we show data from 9 Hopkinson-bar experiments using the material DU along with mean and uncertainty curves. The analysis producing these results assumes a common value for the PTW parameters across all 9 experiments. We refer to this model fitting approach as the “common model analysis”—there is one set of PTW parameters that applies to all experimental conditions. If one looks at the uncertainty from the common model analysis it is clear that the estimated uncertainty in the PTW parameters is not sufficient to cover the observed behavior. Using parameter ranges from this analysis in a dynamic simulation could give misleading results, particularly if the simulation takes the material through states that are not well represented in the data.

The right-hand plot of Figure 1 shows the same data as in the left-hand plot along with mean and uncertainty curves derived from fitting a hierarchical model to the 9 data sets. The hierarchical model allows a separate set of parameters for each experimental sample, it shrinks the parameter estimates together, if possible, and then uses the variation between these estimates to derive parameter uncertainties. The uncertainty curves in Figure 1 are 90% point-wise bounds. This means that for a fixed temperature and strain rate and for a given value of strain, there is a 90% chance that the true stress falls between the corresponding points on the upper and lower uncertainty curves if the assumptions of the model hold.

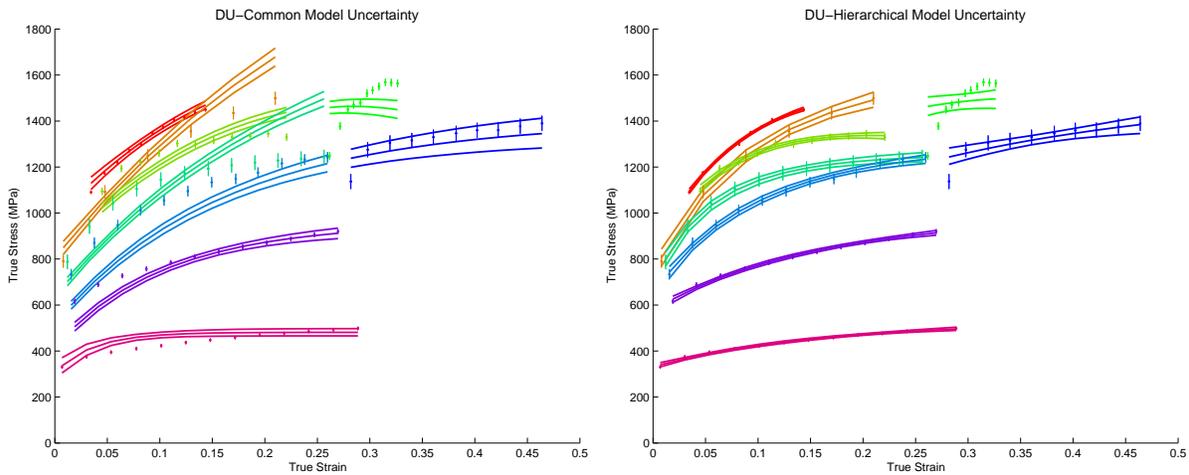


Figure 1: Fitted values and uncertainties for each of the 9 DU Hopkinson-bar experiments. The left-hand plot shows the posterior mean and point-wise 90% uncertainty bounds under the common model. The right-hand plot shows the posterior mean and point-wise 90% uncertainty bounds under the hierarchical model.

In what follows we sometimes refer to data sets that are used to estimate model parameters as “training data” or “training sets.” Data sets that are not used to estimate model parameters are sometimes referred to as “hold-out sets.” To see how the uncertainties from the hierarchical model differ from those of the common model, consider Figure 2. Here predictions were made for a sequence of 18 identical, or replicate, Hopkinson-bar experiments in which 18 samples of DU were subjected to deformation at a temperature of $T = 296K$ and a strain rate of $\dot{\psi} = 0.001/s$. As described in Section 5.3, the common and hierarchical models were each fit to the $K = 9$ DU training sets shown in Figure 1. Based on the fitted model parameters the 18 replicate hold-out sets were then predicted.

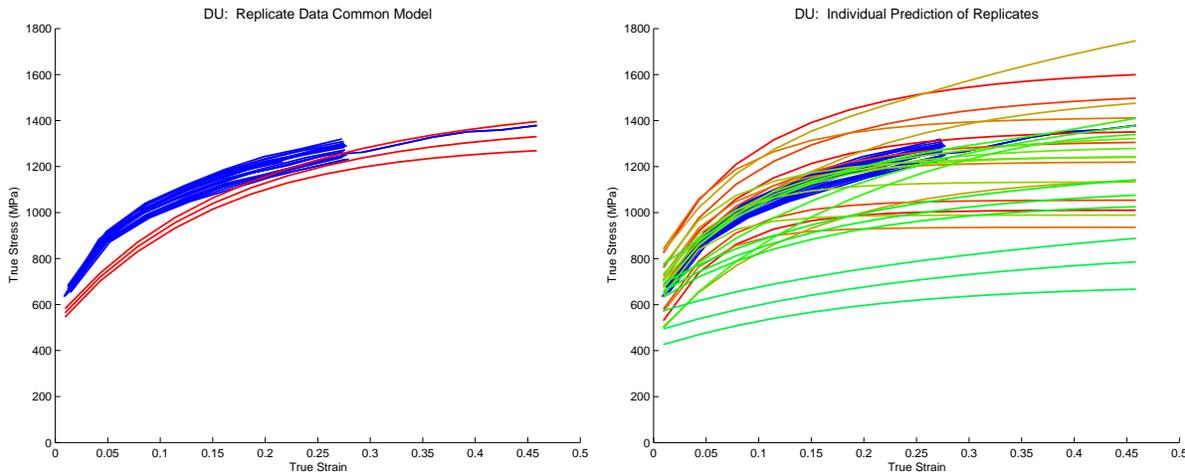


Figure 2: Fitted values and uncertainties for DU Hopkinson-bar experiments at temperature $T = 296K$ and strain rate $\dot{\psi} = 0.001/s$. The thin blue lines show experimental data from 18 different hold-out samples. The left-hand plot shows the posterior mean and point-wise 90% uncertainty bounds under the common model. The right-hand plot shows each of the $K = 9$ posterior mean and point-wise 90% uncertainty bounds under the hierarchical model. The break in lines at the strain value of 0.3 is due to one experiment being stopped and then restarted.

The uncertainty resulting from the common model analysis is shown in the left-hand plot of Figure 2. This uncertainty is clearly too narrow to account for the data. The right-hand plot provides prediction bounds obtained from the hierarchical model fit. The hierarchical model provides separate parameter estimates for each of the 9 DU training sets. Predictions are made for each of the 9 sets of parameter estimates under the temperature and strain rate conditions of the replicate experiments. Thus, in the right-hand plot, there are 9 sets of prediction bounds and mean curves. The overall prediction uncertainty from the hierarchical model can be combined into a single uncertainty as shown in Figure 3.

This uncertainty is generally too conservative for predicting the result of a single Hopkinson-bar experiment. However, our goal is to determine a range of PTW parameter settings that is appropriate over a wide range of temperatures, strains, and strain rates that may be required for a dynamic simulation. Hence this wider uncertainty, resulting from the hierarchical model analysis, is necessary to account for possible material behavior over a wide range of input conditions.

If it is known in advance that a particular process of interest will primarily operate at a particular temperature, strain, and strain rate regime, then experimental data near that particular regime has the potential to reduce uncertainties. If such data is not available,

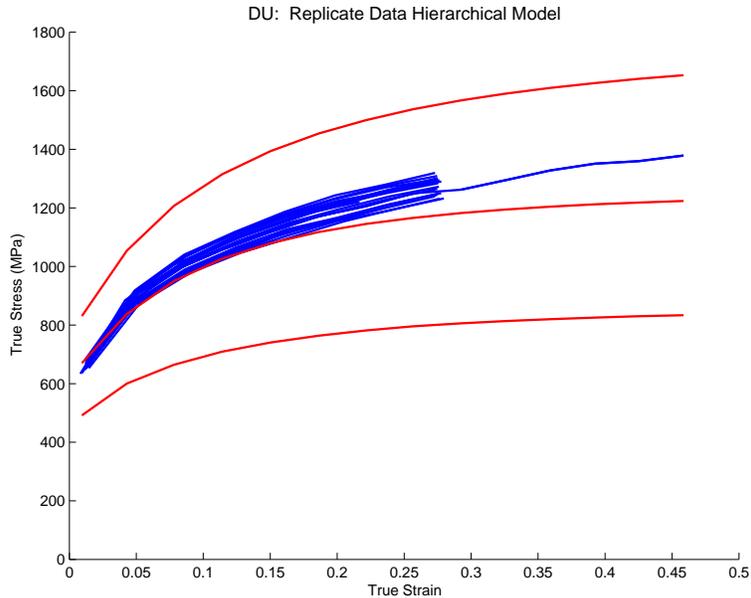


Figure 3: Estimated prediction uncertainties under the hierarchical modeling framework. The hierarchical model uses individual parameter uncertainties to determine combined uncertainty on the PTW parameters, which results in wider prediction intervals. The top and bottom red curves are point-wise 90% uncertainty bands and the middle red curve is the mean.

then one is forced to use the available data to predict behavior in untested regimes. For the Taylor cylinder simulation, one would like the estimated parameter uncertainties to be appropriate for much higher strain rates than used in the Hopkinson-bar experiments. A pressing question is whether the estimated ranges are valid for such extrapolations in strain rate. To see what effect mild extrapolation has on uncertainty estimates we can estimate the PTW model parameters by training on data from low strain rate experiments and then predicting higher strain rate hold-out experiments. Figure 4 shows the results from such an investigation and the results demonstrate there is little additional uncertainty due to this extrapolation. This suggests these estimated ranges are sufficiently broad for extrapolating to high strain rates.

Results from these material strength parameter uncertainty studies have been utilized in several applications involving extrapolation to higher strain rate regimes. For example, flyer plate experiments are conducted to explore the behavior of materials subjected to strain rates several orders of magnitude higher than can be achieved by Hopkinson-bar experiments. Flyer plate experiments involve forcing a plane shock wave through stationary test samples of material and measuring the free surface velocity at the back side of the target as a function of time. This experiment can be modeled as a function of parameters governing the equation of state (EOS), strength and damage properties of the target material. Interest lies in calibrating these parameters to experimental data (e.g. [5]). Uncertainties in the material strength parameters derived from the analysis proposed in this paper provide prior information for this calibration.

Figure 5 shows flyer plate velocity profiles calculated at 500 samples drawn from the final distribution of the PTW parameters for both the common and hierarchical models (EOS and damage parameters were set at nominal values). Note how little uncertainty is present in the traces based on the common model relative to the hierarchical model. This result

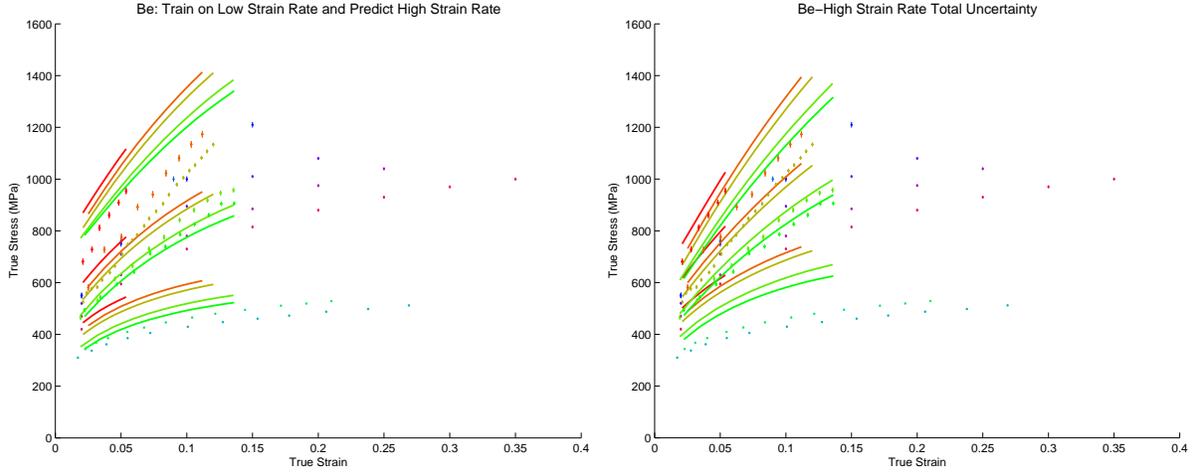


Figure 4: Estimated point-wise 90% prediction uncertainties under the hierarchical modeling framework training on only low strain rate data (left-hand plot) and on the entire data (right-hand plot). The similar uncertainties in the two sets of predictions suggest strain rate extrapolations add little additional uncertainty. In both plots we are predicting the high strain rate data. The prediction bands shown correspond to high strain rate data.

stresses the importance of carefully deriving parameter uncertainties that are designed to apply across a wide range of experimental conditions (such as strain rate and temperature). It will often be the case that assuming a common model for parameters when not required by physical considerations will yield unrealistically small uncertainties, particularly for follow-on analyses involving extrapolations to regimes not represented by the training data. The analysis proposed in this paper maintains the common model as a special case, in the event it is actually the most appropriate representation of the physical process generating the training data.

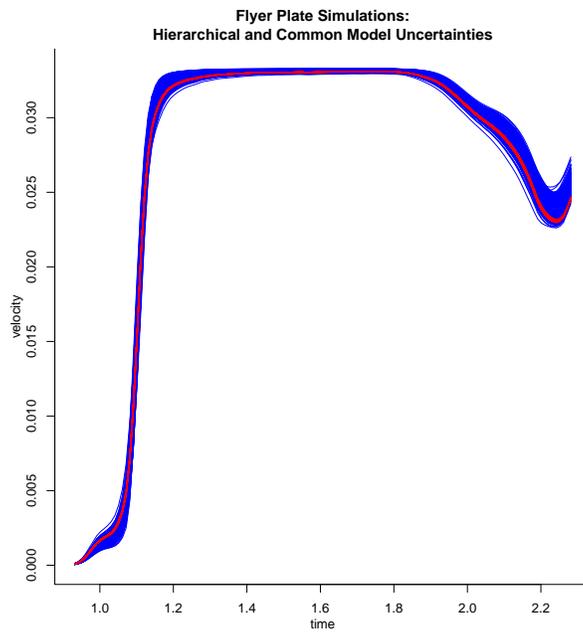


Figure 5: Flyer plate traces corresponding to the final distribution of PTW parameters from the hierarchical model (blue) and common model (red). Each set contains 500 velocity curves. There is very little variability in the common model traces relative to the hierarchical model traces.

4 Bayesian Models and Posterior Distributions

In this section we discuss Bayesian hierarchical models as they apply to modeling data from Hopkinson-bar and quasi-static experiments. We also discuss a special case of a hierarchical model that we refer to as the “common model”. Details about the algorithm we used can be found in the appendix. For more details on Bayesian data analysis and in particular hierarchical models see ([6], [7]).

For each material type a measurement error model for the data can be written as

$$y_{ij} = g(x_{ij}; \mathbf{t}_{ij}, \boldsymbol{\eta}_i) + \epsilon_{ij}, \quad i = 1, \dots, K, \quad j = 1, \dots, n_i, \quad (1)$$

where y_{ij} is the j th observed stress from the i th Hopkinson-bar or quasi-static experiment and $g(x_{ij}; \mathbf{t}_{ij}, \boldsymbol{\eta}_i)$ denotes the PTW model output. The input parameters $\boldsymbol{\eta}_i$ are unknown and must be estimated. The x_{ij} are known calculated strains and \mathbf{t}_{ij} represents the known temperature, strain rate and other fixed parameters corresponding to the data y_{ij} . In what follows we omit x_{ij} and \mathbf{t}_{ij} and write $g_i(\boldsymbol{\eta}_i)$ to denote the PTW model output for data set i . The ϵ_{ij} s represent random measurement errors and we assume they are distributed as independent normal random variables with mean 0 and variance ξ_{ij} . For simplifying formulas we define $\mathbf{y}_i = (y_{i1}, \dots, y_{in_i})^T$; $\mathbf{y} = \text{vec}[\mathbf{y}_1; \dots; \mathbf{y}_K]$, $\boldsymbol{\eta} = \text{vec}[\boldsymbol{\eta}_1; \dots; \boldsymbol{\eta}_K]$, and $\boldsymbol{\xi} = \text{vec}[\boldsymbol{\xi}_1; \dots; \boldsymbol{\xi}_K]$, where the vec operator stacks the column vectors into a single column vector.

Model (1) is a non-linear regression model and the unknown $\boldsymbol{\eta}_i$ parameters can be estimated by independently fitting K regression models. However, estimating the $\boldsymbol{\eta}_i$ s this way ignores information the experiments may share about parameter values. The hierarchical model we propose allows information from each experiment to be linked in the process of estimating the $\boldsymbol{\eta}_i$ s.

One way to combine information from the K experiments is to assume that $\boldsymbol{\eta}_1 = \dots = \boldsymbol{\eta}_K = \boldsymbol{\eta}_0$ and to fit the model

$$y_{ij} = g_i(\boldsymbol{\eta}_0) + \epsilon_{ij}, \quad i = 1, \dots, K, \quad j = 1, \dots, n_i \quad (2)$$

to the observed data. Previous modeling of Hopkinson-bar and quasi-static experiments took this approach, see ([2], [8]). This model is identical to model (1) except that the unknown parameters $\boldsymbol{\eta}_0$ are the same for each experiment. We refer to model (2) as the “common model.” An optimization technique, such as non-linear least squares, is typically used to estimate the unknown parameters $\boldsymbol{\eta}_0$. Uncertainty in the parameter estimates is estimated by appealing to asymptotic approximations (e.g. [9]).

An alternative method for estimating parameters and uncertainties in model (2) is to use a Bayesian modeling approach. A Bayesian treatment of model (2) assumes that the unknown parameters ($\boldsymbol{\eta}_0$ and $\boldsymbol{\xi}$) are random variables and uncertainty about these parameters is described by a prior probability distribution, which we write $p(\boldsymbol{\eta}_0, \boldsymbol{\xi})$. Conditional on $\boldsymbol{\eta}_0$ and $\boldsymbol{\xi}$ the data have some probability of being observed and we denote this by $p(\mathbf{y} | \boldsymbol{\eta}_0, \boldsymbol{\xi})$. Viewed as a function of the parameters ($\boldsymbol{\eta}_0, \boldsymbol{\xi}$), this quantity is referred to as the (data) likelihood function.

The key quantity in a Bayesian analysis of model (2) is the posterior distribution of the unknown parameters, which is

$$p(\boldsymbol{\eta}_0, \boldsymbol{\xi} | \mathbf{y}) = \frac{p(\mathbf{y} | \boldsymbol{\eta}_0, \boldsymbol{\xi}) p(\boldsymbol{\eta}_0, \boldsymbol{\xi})}{p(\mathbf{y})} \propto p(\mathbf{y} | \boldsymbol{\eta}_0, \boldsymbol{\xi}) p(\boldsymbol{\eta}_0, \boldsymbol{\xi}), \quad (3)$$

where $p(\mathbf{y}) = \int p(\mathbf{y} | \boldsymbol{\eta}_0, \boldsymbol{\xi}) p(\boldsymbol{\eta}_0, \boldsymbol{\xi}) d\boldsymbol{\eta}_0 d\boldsymbol{\xi}$ is the marginal distribution of the data which does not depend upon $\boldsymbol{\eta}_0$ or $\boldsymbol{\xi}$. The posterior distribution is proportional to the likelihood function times the prior distribution. The posterior distribution of $\boldsymbol{\eta}_0$, for example, is obtained by integrating $\boldsymbol{\xi}$ out of the left-hand side of (3). From the posterior distribution we can derive interval and point estimates for the unknown parameters and for functions of the parameters.

In the Bayesian framework uncertainty about future observables is described by the posterior predictive distribution. For a future stress value, say \tilde{y} , the posterior predictive distribution is

$$p(\tilde{y} | \mathbf{y}) = \int p(\tilde{y}, \boldsymbol{\theta} | \mathbf{y}) d\boldsymbol{\theta} = \int p(\tilde{y} | \boldsymbol{\theta}, \mathbf{y}) p(\boldsymbol{\theta} | \mathbf{y}) d\boldsymbol{\theta} = \int p(\tilde{y} | \boldsymbol{\theta}) p(\boldsymbol{\theta} | \mathbf{y}) d\boldsymbol{\theta}$$

where for ease of notation we let $\boldsymbol{\theta} = (\boldsymbol{\eta}_0, \boldsymbol{\xi})$. The last line follows because conditional on $\boldsymbol{\theta}$, \tilde{y} and \mathbf{y} are independent. Any quantity of interest, e.g. prediction intervals, about future observables can be obtained from $p(\tilde{y} | \mathbf{y})$.

To implement a Bayesian analysis of model (2) we need to specify a likelihood and a prior distribution. In our work below we assume that $\boldsymbol{\xi}_1 = \dots = \boldsymbol{\xi}_K = \boldsymbol{\xi}_0$ and that $\boldsymbol{\eta}_0$ and $\boldsymbol{\xi}_0$ are independent. The parameters in the PTW model are constrained to lie in some set C (c.f. Table 1) and we incorporate these constraints into the prior distribution for $\boldsymbol{\eta}_0$. For the common model the likelihood and prior, with constraints, are

$$\begin{aligned} p(\mathbf{y}_i | \boldsymbol{\eta}_0, \boldsymbol{\xi}_0) &\sim \mathcal{N}(\mathbf{g}_i(\boldsymbol{\eta}_0), \boldsymbol{\Sigma}_i(\boldsymbol{\xi}_0)), \quad i = 1, \dots, K \\ p(\boldsymbol{\eta}_0, \boldsymbol{\xi}_0) &\propto \mathcal{N}(\mathbf{b}_0, \mathbf{V}_0^{-1}) I(\boldsymbol{\eta}_0 \in C) p(\boldsymbol{\xi}_0). \end{aligned}$$

This notation indicates that conditional on $\boldsymbol{\eta}_0$ and $\boldsymbol{\xi}_0$, the \mathbf{y}_i s have independent normal distributions with mean vector $\mathbf{g}_i(\boldsymbol{\eta}_0)$ and covariance matrix $\boldsymbol{\Sigma}_i(\boldsymbol{\xi}_0)$. The prior distribution for $\boldsymbol{\eta}_0$ is a normal distribution with known mean vector \mathbf{b}_0 and known covariance matrix \mathbf{V}_0^{-1} . To account for the constraints on $\boldsymbol{\eta}_0$ the normal distribution is multiplied by the indicator function $I(\boldsymbol{\eta}_0 \in C)$. The function $I(\boldsymbol{\eta}_0 \in C)$ takes the value 1 if $\boldsymbol{\eta}_0 \in C$ and 0 otherwise. In the applications of this paper, we take $\boldsymbol{\xi}_0 = \lambda$ and $\boldsymbol{\Sigma}_i(\boldsymbol{\xi}_0) = (\sigma_i^2/\lambda) \mathbf{I}_{n_i}$, where σ_i^2 is the measurement error for the i -th experimental data set and \mathbf{I}_{n_i} is the $n_i \times n_i$ identity matrix. The prior distribution for λ is gamma with known shape parameter a and rate parameter b . Because the PTW model is non-linear, we cannot derive an analytical expression for the posterior distribution of $\boldsymbol{\eta}_0$ and $\boldsymbol{\xi}_0$. Therefore, we use Markov-Chain Monte Carlo (MCMC) techniques to obtain samples from the joint posterior distribution, $p(\boldsymbol{\eta}_0, \lambda | \mathbf{y})$, see [6].

In our initial modeling effort, forcing the $\boldsymbol{\eta}_i$ s to be equal and fitting the common model, led to results that appeared inadequate. This led us to consider fitting a hierarchical model. The hierarchical model we propose allows each data set to be fit with a possibly different $\boldsymbol{\eta}_i$ term while still controlling the extent to which the $\boldsymbol{\eta}_i$ s are allowed to differ from each other.

The hierarchical model, with constraints, is:

$$\begin{aligned} p(\mathbf{y}_i | \boldsymbol{\eta}_i, \boldsymbol{\xi}_i) &\sim \mathcal{N}(\mathbf{g}_i(\boldsymbol{\eta}_i), \boldsymbol{\Sigma}_i(\boldsymbol{\xi}_i)), \quad i = 1, \dots, K \\ p(\boldsymbol{\eta}_i | \boldsymbol{\eta}_0, \mathbf{V}_\boldsymbol{\eta}) &\propto \mathcal{N}(\boldsymbol{\eta}_0, \mathbf{V}_\boldsymbol{\eta}^{-1}) I(\boldsymbol{\eta}_i \in C) \\ p(\boldsymbol{\eta}_0) &\propto \mathcal{N}(\mathbf{b}_0, \mathbf{V}_0^{-1}) I(\boldsymbol{\eta}_0 \in C) \\ p(\mathbf{V}_\boldsymbol{\eta}) &\sim W((\nu\boldsymbol{\Omega})^{-1}, \nu). \end{aligned} \tag{4}$$

This notation indicates that conditional on the random variables $\boldsymbol{\eta}_i$ and $\boldsymbol{\xi}_i$, the \mathbf{y}_i s have independent normal distributions with mean vectors $\mathbf{g}_i(\boldsymbol{\eta}_i)$ and covariance matrices $\boldsymbol{\Sigma}_i(\boldsymbol{\xi}_i)$.

Similarly, conditional on $\boldsymbol{\eta}_0$ and \mathbf{V}_η , the $\boldsymbol{\eta}_i$ s have independent normal distributions with mean vector $\boldsymbol{\eta}_0$ and covariance matrix \mathbf{V}_η^{-1} , subject to the constraints $\boldsymbol{\eta}_i \in C$, $i = 1, \dots, K$. The extent to which the $\boldsymbol{\eta}_i$ s differ from each other will depend on the structure of the covariance matrix \mathbf{V}_η^{-1} . It is frequently more convenient to work with precision matrices instead of covariance matrices. A precision matrix is simply the inverse of a covariance matrix; for example, the precision matrix for $\boldsymbol{\eta}_i$ is \mathbf{V}_η .

The prior distribution for $\boldsymbol{\eta}_0$ is normal with known mean vector \mathbf{b}_0 and known covariance matrix \mathbf{V}_0^{-1} , subject to the constraint $\boldsymbol{\eta}_0 \in C$. The prior distribution for \mathbf{V}_η is the Wishart distribution with positive definite scale matrix $\boldsymbol{\Omega}$ and ν degrees of freedom. The quantities \mathbf{b}_0 , \mathbf{V}_0 , $\boldsymbol{\Omega}$ and ν are all specified and discussed in more detail in the appendix. In the discussion that follows we will not explicitly write down the constraint conditions but the reader should keep in mind that the support of *all* conditional distributions involving the $\boldsymbol{\eta}_i$ s and $\boldsymbol{\eta}_0$ depend on C . The prior distributions specified in lines 2–4 of (4) are chosen mainly for mathematical convenience in that they facilitate a Gibbs sampling MCMC implementation as described below and in the appendix. There is flexibility to implement a wide variety of prior beliefs through proper choice of \mathbf{b}_0 , \mathbf{V}_0 , $\boldsymbol{\Omega}$ and ν . Closely related model specifications have been examined in the statistical literature (e.g. [7] and the references therein).

The posterior distribution of all the parameters given the observed data is,

$$p(\boldsymbol{\eta}, \boldsymbol{\xi}, \boldsymbol{\eta}_0, \mathbf{V}_\eta | \mathbf{y}) \propto p(\mathbf{y} | \boldsymbol{\eta}, \boldsymbol{\xi}) p(\boldsymbol{\eta} | \boldsymbol{\eta}_0, \mathbf{V}_\eta) p(\boldsymbol{\xi}) p(\boldsymbol{\eta}_0) p(\mathbf{V}_\eta).$$

This formulation assumes that $\boldsymbol{\eta}_0$ and \mathbf{V}_η are independent and that $\boldsymbol{\xi}$ is independent of $\boldsymbol{\eta}$, $\boldsymbol{\eta}_0$, and \mathbf{V}_η . Because this posterior distribution is not analytically tractable we use the Gibbs sampler ([10]) to generate samples from this distribution. With samples from the posterior distribution we can determine interval and point estimates for quantities of interest. The Gibbs sampler requires the full conditional distribution for each parameter. The full conditionals are:

$$\begin{aligned} p(\boldsymbol{\eta} | \mathbf{y}, \boldsymbol{\xi}, \boldsymbol{\eta}_0, \mathbf{V}_\eta) &\propto p(\mathbf{y} | \boldsymbol{\eta}, \boldsymbol{\xi}) p(\boldsymbol{\eta} | \boldsymbol{\eta}_0, \mathbf{V}_\eta) \\ p(\boldsymbol{\xi} | \mathbf{y}, \boldsymbol{\eta}, \boldsymbol{\eta}_0, \mathbf{V}_\eta) &\propto p(\mathbf{y} | \boldsymbol{\eta}, \boldsymbol{\xi}) p(\boldsymbol{\xi}) \\ p(\boldsymbol{\eta}_0 | \mathbf{y}, \boldsymbol{\eta}, \boldsymbol{\xi}, \mathbf{V}_\eta) &\propto p(\boldsymbol{\eta} | \boldsymbol{\eta}_0, \mathbf{V}_\eta) p(\boldsymbol{\eta}_0) \\ p(\mathbf{V}_\eta | \mathbf{y}, \boldsymbol{\eta}, \boldsymbol{\xi}, \boldsymbol{\eta}_0) &\propto p(\boldsymbol{\eta} | \boldsymbol{\eta}_0, \mathbf{V}_\eta) p(\mathbf{V}_\eta). \end{aligned} \tag{5}$$

Explicit expressions for these components are derived in the appendix. The Gibbs algorithm provides samples from the joint posterior distribution that are used for uncertainty quantification.

Inference for individual $\boldsymbol{\eta}_i$ parameters is based on the posterior distribution $p(\boldsymbol{\eta} | \mathbf{y})$. This distribution conveys the uncertainty in the parameter estimates associated with the K experiments used to fit the hierarchical model.

Let $\boldsymbol{\eta}_{K+1}$ represent the PTW model parameters corresponding to experimental conditions covering the strain, strain rate, and temperature range seen in the data used to estimate the hierarchical model. In general, there will be more uncertainty in the estimate of $\boldsymbol{\eta}_{K+1}$ than in any of the individual $\boldsymbol{\eta}_i$ s. Inferences about $\boldsymbol{\eta}_{K+1}$ are based on the posterior distribution

$p(\boldsymbol{\eta}_{K+1} | \mathbf{y})$, which is

$$\begin{aligned} p(\boldsymbol{\eta}_{K+1} | \mathbf{y}) &= \int p(\boldsymbol{\eta}_{K+1} | \boldsymbol{\eta}_0, \mathbf{V}_\eta) p(\boldsymbol{\eta}_0, \mathbf{V}_\eta | \mathbf{y}) d\boldsymbol{\eta}_0 d\mathbf{V}_\eta \\ &\approx \frac{1}{N} \sum_{j=1}^N p(\boldsymbol{\eta}_{K+1} | \boldsymbol{\eta}_{0,j}, \mathbf{V}_{\eta,j}). \end{aligned} \quad (6)$$

The Gibbs algorithm provides realizations $\boldsymbol{\eta}_{0,j}$ and $\mathbf{V}_{\eta,j}$ from $p(\boldsymbol{\eta}_0, \mathbf{V}_\eta | \mathbf{y})$ and the integral is approximated with a sample average over N MCMC realizations.

Inference about $\boldsymbol{\eta}_0$ is made through the posterior distribution $p(\boldsymbol{\eta}_0 | \mathbf{y})$, which is

$$\begin{aligned} p(\boldsymbol{\eta}_0 | \mathbf{y}) &= \int p(\boldsymbol{\eta}_0 | \mathbf{y}, \boldsymbol{\eta}, \mathbf{V}_\eta) p(\boldsymbol{\eta}, \mathbf{V}_\eta | \mathbf{y}) d\boldsymbol{\eta} d\mathbf{V}_\eta \\ &\approx \frac{1}{N} \sum_{j=1}^N p(\boldsymbol{\eta}_0 | \mathbf{y}, \boldsymbol{\eta}_j, \mathbf{V}_{\eta,j}) \end{aligned} \quad (7)$$

where $\boldsymbol{\eta}_j$ and $\mathbf{V}_{\eta,j}$ are samples from $p(\boldsymbol{\eta}, \mathbf{V}_\eta | \mathbf{y})$. In the appendix we show that $p(\boldsymbol{\eta}_0 | \mathbf{y}, \boldsymbol{\eta}, \mathbf{V}_\eta)$ is equal to the full conditional distribution $p(\boldsymbol{\eta}_0 | \mathbf{y}, \boldsymbol{\eta}, \mathbf{V}_\eta, \boldsymbol{\xi})$, facilitating computation of (7).

Posterior predictive samples of PTW model output are generated by evaluating the PTW model at specified strains \mathbf{x} for each of N posterior samples $\boldsymbol{\eta}_k^P$. Specified quantiles of these curves at strain value x_j are computed from the N realizations $\{g(x_j, \boldsymbol{\eta}_1^P), \dots, g(x_j, \boldsymbol{\eta}_N^P)\}$. For example, the curves in the right-hand panels of Figures 1 and 2 were generated for data set i by setting $\boldsymbol{\eta}_k^P = \boldsymbol{\eta}_{i,k}$, where $\boldsymbol{\eta}_{i,k}$ is the k th posterior sample of $\boldsymbol{\eta}_i$, the PTW model parameters corresponding to data set i . The curves in Figure 3 and in both panels of Figure 4 are derived from setting $\boldsymbol{\eta}_k^P = \boldsymbol{\eta}_{K+1,k}$, where $\boldsymbol{\eta}_{K+1,k}$ is the k th posterior sample generated according to (6).

The hierarchical model of equation (4) can be reduced to the common parameter model (2) by taking $\boldsymbol{\xi}_1 = \dots = \boldsymbol{\xi}_K = \boldsymbol{\xi}_0$ and by letting $\boldsymbol{\Omega} \rightarrow \mathbf{0}$ and $\nu \rightarrow \infty$ in equation (4). These limits in $\boldsymbol{\Omega}$ and ν result in precision matrices \mathbf{V}_η that diverge to infinity, or covariance matrices that converge to the zero matrix, so that the $\boldsymbol{\eta}_i$ s are degenerate random variables that equal their mean value of $\boldsymbol{\eta}_0$. This suggests that the common model is adequate if the posterior distributions for $\boldsymbol{\eta}_{K+1}$ and $\boldsymbol{\eta}_0$ are nearly the same.

5 Analysis of Material Characterization Experiments

We now outline the analysis of the data from the material characterization experiments. Results from the following materials are presented in this section: aluminum (Al), beryllium (Be), depleted uranium (DU), tantalum (Ta), and uranium 6 wt% niobium (U-6Nb).

Basic stress-strain data at moderate strain rates (about 10^3 s^{-1}) are typically obtained in a Hopkinson-bar experiment in which an elastic wave is passed through a thin cylinder of the material under investigation. Strain gauges mounted on the support cylinders measure strain as a function of time. From these measurements, the stress-strain behavior of the material is inferred. The data from a Hopkinson-bar experiment on tantalum done at room temperature (298K) and a strain rate of 1300/s are shown in Figure 6.

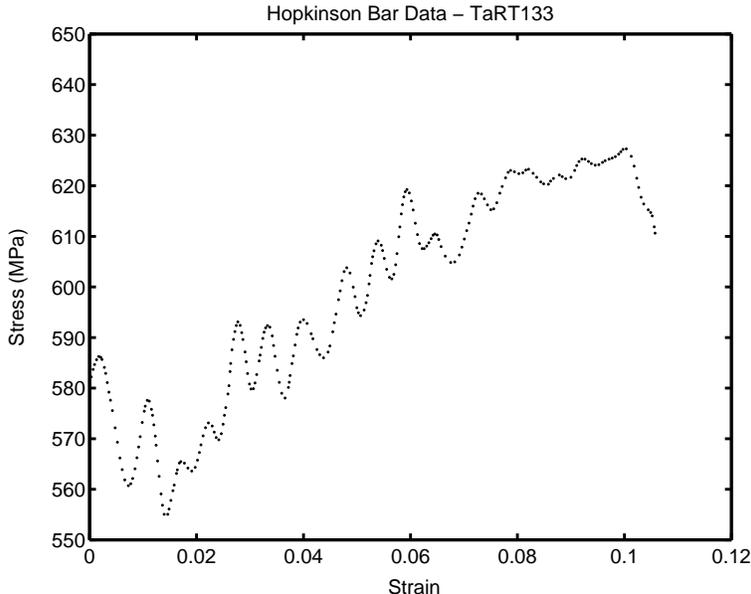


Figure 6: Plot of data obtained from Hopkinson-bar experiment done on tantalum at room temperature and a strain rate of 1300 s^{-1} .

This figure shows a well-known feature of Hopkinson-bar experiments, the presence of wiggles in the measured stress as a function of strain, which are particularly evident at strains of 0.02 and below but also observable at higher strains. These oscillations, caused by elastic wave dispersion within the sample and apparatus, tend to reduce the accuracy of Hopkinson-bar data. The data below strains of 0.017 seem unreliable because of their higher amplitude of oscillations and the fact that the stress rises as the strain approaches zero instead of falling. Likewise, the data above a strain of 0.1 seem to be corrupted by an artifact. In the present analysis, the data in these two end regions are excluded. To make use of the data between 0.017 and 0.1, their uncertainties need to be quantitatively characterized.

For the materials beryllium and tantalum, we refer the reader to [8] for a explanation of how the data was thinned before our analysis. For each data set of the remaining materials, we sub-sampled observations at constant frequency throughout the data set and fit a quadratic function to the thinned data, computed the residuals and then examined the autocorrelation of the residuals. The number of points to thin was chosen so that the autocorrelation in the residuals was nearly 0.

Tables 3, 4, and 5 give a brief description of the PTW parameters to estimate, the fixed PTW parameters, and physical constants, respectively. In each subsection we provide the values used in model fitting.

Parameter	Description
θ	Initial strain hardening rate
κ	Constant relates to temperature dependence
$-\log(\gamma)$	Constant relates to strain dependence
y_0	Maximum yield stress (at 0K)
y_∞	Minimum yield stress (melting)
s_0	Maximum saturation stress (at 0K)
s_∞	Minimum saturation stress (melting)

Table 3: PTW parameters to estimate. All parameters are dimensionless.

Parameter	Description
p	Constant modifying Voce hardening law
y_1	Constant in high strain rate regime
y_2	Constant in high strain rate regime
β	Constant in high strain rate regime

Table 4: Fixed PTW parameters. All parameters are dimensionless.

Parameter	Description	Units
T_m	Melting temperature	K
G_0	Shear modulus (at 0K)	kilobars
α	Constant in $G(T)$ equation	dimensionless
C_v	Heat Capacity	J/Kg/K
A	Atomic weight	dalton
ρ	Density	g/cm^3

Table 5: Empirically determined constants or physical constants.

5.1 Aluminum

In this section we discuss fitting the PTW model to data from Hopkinson-bar experiments on aluminum. There are four aluminum data sets. Table 6 lists the data sets used along with some of the experimental conditions for each data set. In the table, n_i is the number of observations after thinning and σ_i is the root mean square error (RMSE) from fitting a quadratic model to the thinned data.

Data Set	n_i	Temp (K)	Strain Rate (s^{-1})	σ_i
A6.077.1e-3.p.txt	16	77	0.001	10.0
A6.298.2500.p.txt	19	298	2500	6.0
A6.298.1e-1.p.txt	13	298	0.1	6.0
A6.298.1e-3.p.txt	12	298	0.001	4.0

Table 6: Al: Experimental conditions.

Table 7 gives the values we used for the fixed parameters in the PTW model. These values are from [2].

Parameter	Nominal Value	Parameter	Nominal Value
p	3.0	G_0	299.2
y_1	0.0142	α	0.475
y_2	0.40	C_v	898.7
β	0.23	A	27.0836
T_m	932	ρ	2.710

Table 7: Al: Fixed PTW parameters.

Table 8 gives the values we used for \mathbf{b}_0 and \mathbf{V}_0 . The values of \mathbf{b}_0 are from [2]. This table also gives the prior values we used for $\mathbf{\Omega}$. The appendix describes the method used to determine these values.

Parameter	\mathbf{b}_0	$\mathbf{V}_{0,ii}$	$\mathbf{\Omega}_{ii}$
θ	0.0184	1069	1.713×10^{-5}
κ	0.2	14	0.002736
$-\log(\gamma)$	9.903	0.10055	2.658
y_0	0.00898	52488	3.222×10^{-6}
y_∞	0.00596	197531	1.587×10^{-8}
s_0	0.0142	4221	5.172×10^{-5}
s_∞	0.00856	84061	4.347×10^{-8}

Table 8: Al: Values for $\boldsymbol{\eta}_0$ and $\mathbf{V}_\boldsymbol{\eta}$ prior distributions.

Table 9 shows the posterior mean and standard deviation for each parameter and each data set. These were calculated from the MCMC output by taking every other observation

after discarding the first 100,000 samples. Therefore, these estimates are based on 250,000 samples.

Data Set	θ	κ	$-\log(\gamma)$	y_0	y_∞	s_0	s_∞
A6.077.1e-3.p.txt 77K 0.001/s	0.0529 (0.00676)	0.494 (0.124)	8.79 (4.61)	0.00942 (0.00304)	0.00566 (0.000237)	0.032 (0.0125)	0.00791 (0.00055)
A6.298.2500.p.txt 298K 2500/s	0.0254 (0.00391)	0.362 (0.124)	11.6 (3.65)	0.0129 (0.00281)	0.0059 (0.000236)	0.0285 (0.00994)	0.00777 (0.000391)
A6.298.1e-1.p.txt 298K 0.1/s	0.0364 (0.00461)	0.417 (0.119)	10.3 (3.66)	0.0112 (0.00293)	0.00581 (7.84×10^{-5})	0.028 (0.011)	0.00785 (0.000121)
A6.298.1e-3.p.txt 298K 0.001/s	0.0331 (0.00411)	0.41 (0.119)	10.5 (3.66)	0.0114 (0.00302)	0.0058 (6.37×10^{-5})	0.0275 (0.0113)	0.0077 (7.85×10^{-5})

Table 9: Al: Posterior mean (standard deviation).

Figure 7 is a plot of the thinned aluminum data along with $\pm 1\sigma$ error bars.

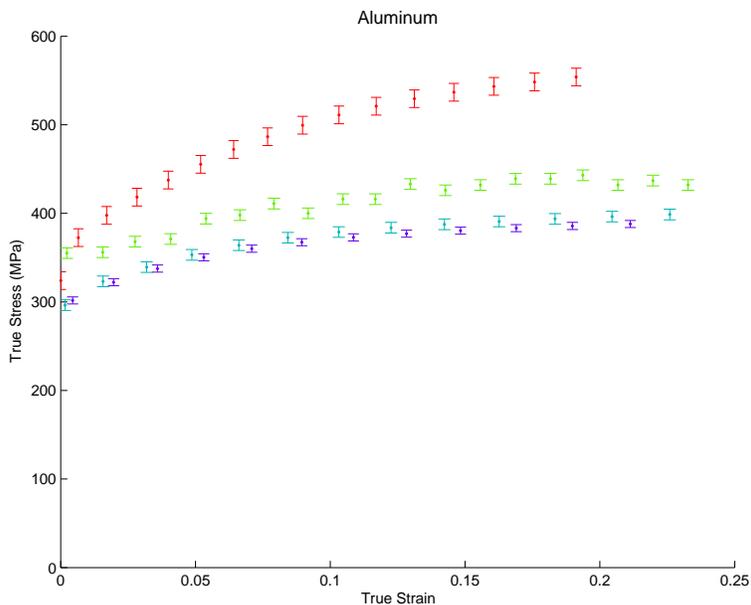


Figure 7: Al: Sampled data with $\pm 1\sigma$ error bars.

Figure 8 shows point-wise posterior prediction intervals derived from fitting the common model and the hierarchical model. In both plots these intervals correspond to the 0.05, 0.50, and 0.95 quantiles of the MCMC output. The common and hierarchical models provide comparable fits to the individual data sets.

Figure 9 shows approximate 95% posterior probability regions. The blue ellipses correspond to the individual data sets. The broken black ellipses correspond to $\boldsymbol{\eta}_0$ and the solid black ellipses correspond to $\boldsymbol{\eta}_{K+1}$. The cyan ellipses are derived from fitting the common model to all the data. Fitting the common model gives a somewhat optimistic picture of our knowledge about the parameters. Based on these plots, fitting a single nonlinear regression model to all the experimental conditions under-represents the amount of uncertainty present in the PTW parameters. For the plots in the lower triangle the points labeled with a “c” are point estimates given in [2]. Note that this plot was constructed by assuming bivariate normality for the marginal posterior distributions, not accounting for constraints on the individual parameters. Therefore, this plot should be viewed merely as a heuristic for assessing

bivariate relationships among the parameter distributions represented.

Figure 10 shows 0.05, 0.50, and 0.95 point-wise posterior prediction intervals based on $\boldsymbol{\eta}_{K+1}$ and $\boldsymbol{\eta}_0$ from the hierarchical model. This plot shows how additional posterior un-

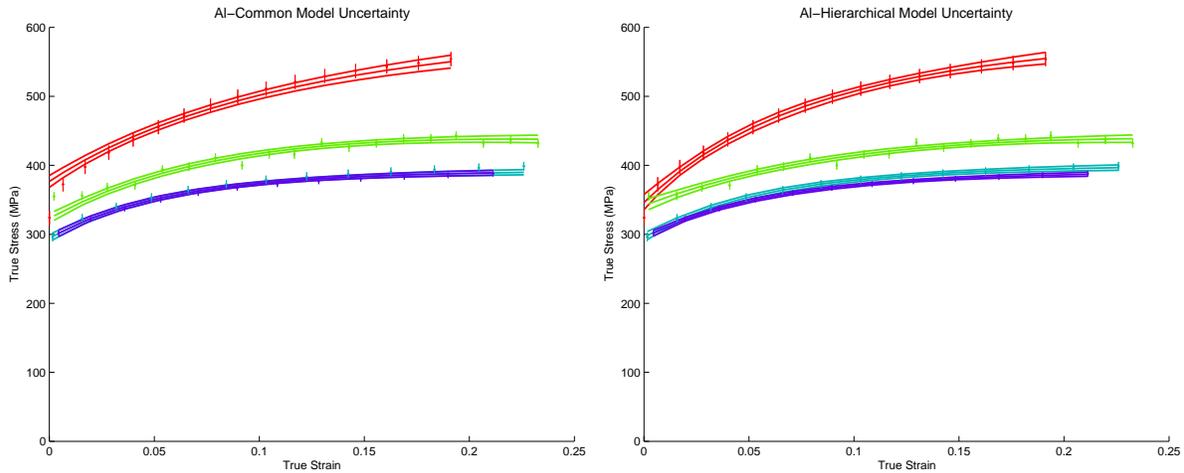


Figure 8: AI: Point-wise posterior prediction intervals from the common model fit, left-hand plot, and the hierarchical model fit, right-hand plot. The curves are 0.05, 0.50, and 0.95 point-wise quantiles from the MCMC output. The curves in the right-hand plot are based on the $\boldsymbol{\eta}_i$ s.

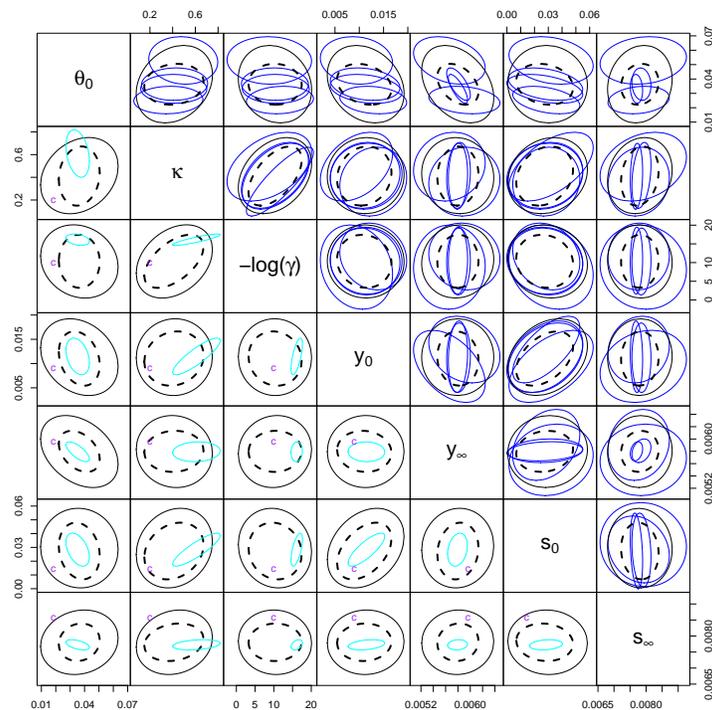


Figure 9: AI: Approximate 95% posterior probability regions. The blue ellipses in the upper right correspond to the $\boldsymbol{\eta}_i$ s, one for each data set. The broken black ellipses and the solid black ellipses correspond to $\boldsymbol{\eta}_0$ and $\boldsymbol{\eta}_{K+1}$, respectively, from the hierarchical model fit. The cyan ellipses in the lower left correspond to $\boldsymbol{\eta}_0$ from the common model fit.

certainty in $\boldsymbol{\eta}_{K+1}$ relative to $\boldsymbol{\eta}_0$ translates into increased posterior prediction uncertainty.

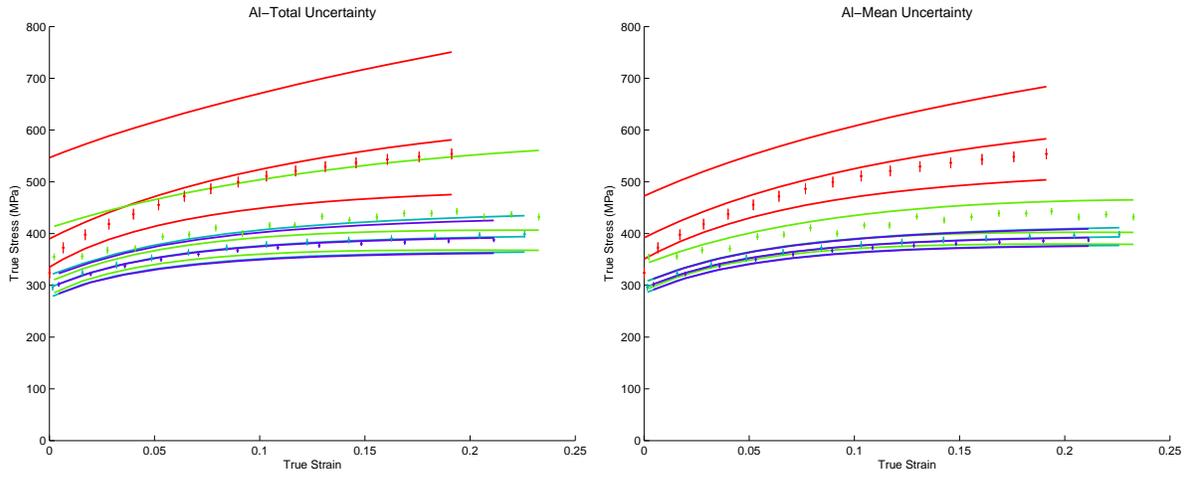


Figure 10: AI: Posterior prediction intervals from the hierarchical model. The left-hand plot is based on $\boldsymbol{\eta}_{K+1}$ and the right-hand plot is based on $\boldsymbol{\eta}_0$. The curves are 0.05, 0.50, and 0.95 point-wise quantiles from the MCMC output.

5.2 Beryllium

In this section we discuss fitting the PTW model to data from Hopkinson-bar and quasi-static experiments on beryllium. There are twelve beryllium data sets.

Table 10 lists the data sets used along with some of the experimental conditions for each data set. In the table, n_i is the number of observations after thinning and σ_i is the root mean square error (RMSE) obtained from applying the fitting method of [8] to the thinned data.

Data Set	n_i	Temp (K)	Strain Rate (s^{-1})	σ_i
BeHP-196303.txt	6	77	3000	11.0
Bed-50353.txt	9	223	3500	12.0
Bef0203500.x.txt	23	298	3500	6.4
Bed20373.txt	11	473	3700	8.9
Bed30393.txt	11	573	3900	7.0
Bef300c1.txt	12	573	1.0	2.3
Bef200c1e-3.txt	12	473	0.001	2.1
BeMon020203.txt	3	293	2000	10.0
BeMon020952.txt	4	293	950	10.0
BeMon020200.txt	5	293	2.0	5.0
BeMon02020-2.txt	6	293	0.02	5.0
BeMon02020-4.txt	8	293	0.0001	5.0

Table 10: Beryllium data sets.

Table 11 lists the values we used for the fixed PTW parameters.

Parameter	Nominal Value	Parameter	Nominal Value
p	2.0	G_0	1524
y_1	0.0077	α	0.32
y_2	0.40	C_v	1820
β	0.25	A	9.013
T_m	1560	ρ	1.85

Table 11: Be: Fixed PTW parameters.

Table 12 gives the values we used for \mathbf{b}_0 and \mathbf{V}_0 . This table also gives the prior values we used for $\mathbf{\Omega}$. The appendix describes the method used to determine these values.

Table 13 gives posterior means and standard deviations for each parameter and data set. These were calculated from the MCMC output by taking every other realization after discarding the first 100,000 samples. Therefore, these estimates are based on 250,000 samples.

Figure 11 is a plot of the data used in fitting along with $\pm 1\sigma$ error bars.

Figure 12 shows point-wise posterior prediction intervals derived from fitting the common model and the hierarchical model. In both plots these intervals correspond to the 0.05, 0.50,

Parameter	b_0	$V_{0,ii}$	Ω_{ii}
θ	0.0394	0.0	4.341×10^{-6}
κ	0.145	0.0	0.000426
$-\log(\gamma)$	13.8155	0.0	2.6461
y_0	0.0018	0.0	1.201×10^{-8}
y_∞	0.0004	0.0	1.846×10^{-8}
s_0	0.0077	0.0	4.014×10^{-6}
s_∞	0.0006	0.0	3.44×10^{-7}

Table 12: Be: Values for η_0 and V_η prior distributions.

Data Set	θ	κ	$-\log(\gamma)$	y_0	y_∞	s_0	s_∞
BeHP-196303.txt 77K 3000/s	0.0357 (0.00346)	0.218 (0.117)	9.33 (6.09)	0.00172 (0.000127)	0.00103 (0.000315)	0.0231 (0.0099)	0.00396 (0.00157)
Bed-50353.txt 223K 3500/s	0.0309 (0.00266)	0.203 (0.112)	10.1 (5.27)	0.00157 (0.000168)	0.000943 (0.000266)	0.0224 (0.00752)	0.00378 (0.00157)
Bef0203500.x.txt 298K 3500/s	0.0259 (0.00101)	0.164 (0.12)	12.0 (5.83)	0.00144 (0.000169)	0.000893 (0.000286)	0.0303 (0.00737)	0.00408 (0.00223)
Bed20373.txt 473K 3700/s	0.0255 (0.00218)	0.196 (0.0996)	11.3 (4.28)	0.00162 (0.000208)	0.000922 (0.000226)	0.0156 (0.00631)	0.00322 (0.00108)
Bed30393.txt 573K 3900/s	0.0207 (0.00206)	0.171 (0.106)	13.4 (4.52)	0.00166 (0.000191)	0.000973 (0.00028)	0.015 (0.00533)	0.00315 (0.00123)
Bef300c1.txt 573K 1.0/s	0.0154 (0.00169)	0.18 (0.11)	13.7 (6.28)	0.00164 (0.000327)	0.000825 (0.000198)	0.00259 (0.000579)	0.00202 (0.000284)
Bef200c1e-3.txt 473K 0.001/s	0.0132 (0.00149)	0.172 (0.117)	14.2 (6.43)	0.00162 (0.000349)	0.000771 (0.000216)	0.00221 (0.000386)	0.00192 (0.000188)
BeMon020203.txt 293K 2000/s	0.0296 (0.00275)	0.199 (0.11)	10.8 (5.12)	0.00164 (0.000177)	0.000978 (0.000264)	0.0208 (0.00784)	0.00365 (0.00141)
BeMon020952.txt 293K 950/s	0.0275 (0.00247)	0.195 (0.105)	11.5 (4.78)	0.00169 (0.000172)	0.000991 (0.000271)	0.0167 (0.0058)	0.00341 (0.0012)
BeMon020200.txt 293K 2.0/s	0.0356 (0.00401)	0.247 (0.104)	8.77 (5.84)	0.00174 (0.000259)	0.00103 (0.000242)	0.0138 (0.00955)	0.00295 (0.000831)
BeMon02020-2.txt 293K 0.02/s	0.0244 (0.00252)	0.19 (0.0964)	12.1 (5.48)	0.00169 (0.000238)	0.000969 (0.000245)	0.00904 (0.00532)	0.00307 (0.000848)
BeMon02020-4.txt 293K 0.0001/s	0.0271 (0.00422)	0.207 (0.103)	11.6 (5.67)	0.00168 (0.000268)	0.000884 (0.000207)	0.00776 (0.005)	0.00314 (0.000641)

Table 13: Be: Posterior mean (standard deviation).

and 0.95 quantiles of the MCMC output. The hierarchical model provides a much better fit to the individual data sets than does the common model.

Figure 13 shows approximate 95% posterior probability regions. The blue and red ellipses correspond to the individual data sets; the blue ellipses are from strain rates below 3000/s and the red ellipses are from strain rates at or above 3000/s. The broken black ellipses correspond to η_0 and the solid black ellipses correspond to η_{K+1} . The cyan ellipses are derived from fitting a common model to all the data. Fitting the common model gives a very optimistic picture of our knowledge about the parameters. Based on these plots, fitting a single nonlinear regression model to all the experimental conditions under-represents the amount of uncertainty present in the PTW parameters. For the plots in the lower triangle, the points labeled with a “p” are point estimates given in [1], “c” are point estimates given in [2], and “b” are point estimates given by Blumenthal and reported in [2]. This plot was constructed by assuming bivariate normality for the marginal posterior distributions, not accounting for constraints on the individual parameters. Therefore, this plot should

be viewed merely as a heuristic for assessing bivariate relationships among the parameter distributions represented.

Figure 14 shows 0.05, 0.50, and 0.95 point-wise posterior prediction intervals based on $\boldsymbol{\eta}_{K+1}$ and $\boldsymbol{\eta}_0$ from the hierarchical model. This plot shows how additional posterior uncertainty in $\boldsymbol{\eta}_{K+1}$ relative to $\boldsymbol{\eta}_0$ translates into increased posterior prediction uncertainty.

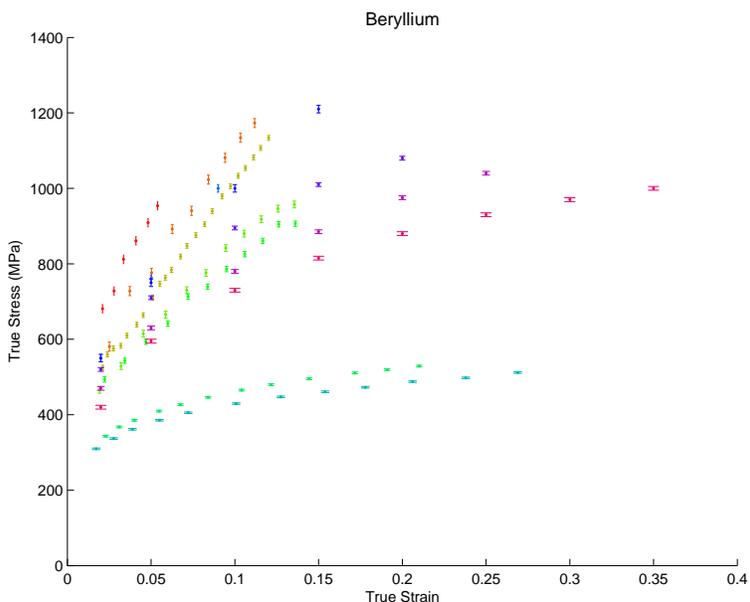


Figure 11: Be: Sampled data with $\pm 1\sigma$ error bars.

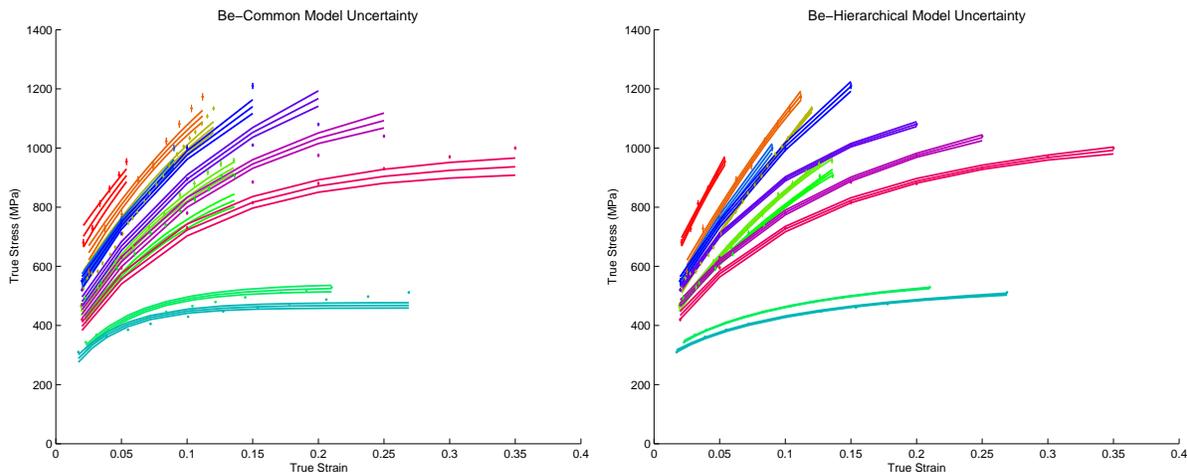


Figure 12: Be: Point-wise posterior prediction intervals from the common model fit, left-hand plot, and the hierarchical model fit, right-hand plot. The curves are 0.05, 0.50, and 0.95 point-wise quantiles from the MCMC output. The curves in the right-hand plot are based on the $\boldsymbol{\eta}_i$ s.

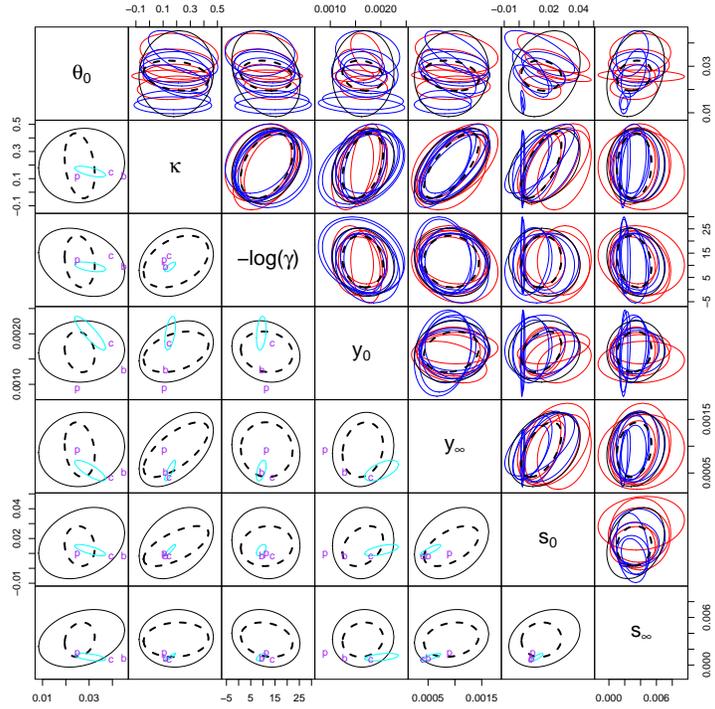


Figure 13: Be: Approximate 95% posterior probability regions. The blue and red ellipses in the upper right correspond to the η_i s, one for each data set; the blue ellipses are from data sets with strain rates below 3000/s and the red ellipses are from data sets with strain rates above 3000/s. The broken black ellipses and the solid black ellipses correspond to η_0 and η_{K+1} , respectively, from the hierarchical model fit. The cyan ellipses in the lower left correspond to η_0 from the common model fit.

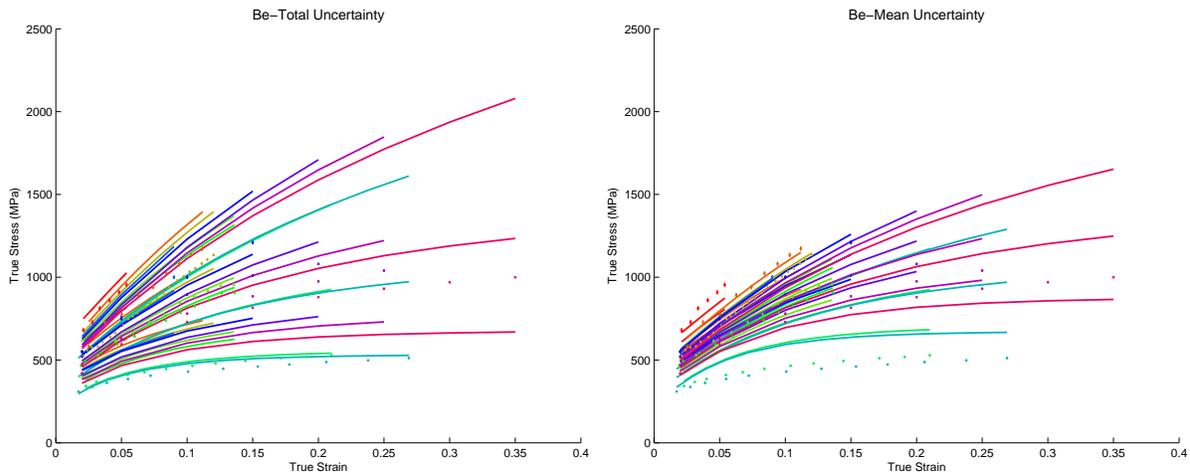


Figure 14: Be: Posterior prediction intervals from the hierarchical model. The left-hand plot is based on η_{K+1} and the right-hand plot is based on η_0 . The curves are 0.05, 0.50, and 0.95 point-wise quantiles from the MCMC output.

5.3 Depleted Uranium

In this section we discuss fitting the PTW model to data from Hopkinson-bar experiments on depleted uranium. There are nine depleted uranium data sets.

Table 14 lists the data sets used along with some of the experimental conditions for each data set. In the table, n_i is the number of observations after thinning and σ_i is the root mean square error (RMSE) from fitting a quadratic model to the thinned data.

Data Set	n_i	Temp (K)	Strain Rate (s^{-1})	σ_i
DUtt-67223.p.txt	11	206	2200	7.0
DUtt-7510-1.p.txt	6	198	0.1	27.0
DUtt023343.p.txt	11	296	3400	15.0
DUtt023283r.p.txt	11	296	2800	16.0
DUtt023100.p.txt	12	296	1.0	31.0
DUtt02310-3b.p.txt	12	296	0.001	22.0
DUtt02310-3br.p.txt	10	296	0.001	32.0
DUtt15010-1.p.txt	12	423	0.1	10.0
DUtt30010-1.p.txt	13	573	0.1	7.0

Table 14: Depleted uranium data sets.

Table 15 gives the values we used for the fixed parameters in the PTW model. These values are from [2].

Parameter	Nominal Value	Parameter	Nominal Value
p	1.0	G_0	938.9
y_1	0.03	α	0.56
y_2	0.27	C_v	118
β	0.27	A	238.1494
T_m	1405	ρ	19.05

Table 15: DU: Fixed PTW parameters.

Table 16 gives the values we used for \mathbf{b}_0 and \mathbf{V}_0 . The values of \mathbf{b}_0 are from [2]. This table also gives the prior values we used for $\mathbf{\Omega}$. The appendix describes the method used to determine these values.

Table 17 gives posterior means and standard deviations for each variable and data set. These were calculated from the MCMC output by taking every other realization after discarding the first 100,000 samples. Therefore, these estimates are based on 250,000 samples.

Figure 15 is a plot of the observed data from one of the Hopkinson-bar experiments. The data for all of the experimental conditions along with $\pm 1\sigma$ error bars is shown in Figure 16.

Figure 17 shows point-wise posterior prediction intervals derived from fitting the common model and the hierarchical model. In both plots these intervals correspond to the 0.05, 0.50,

Parameter	b_0	$V_{0,ii}$	Ω_{ii}
θ	0.088	1069	3.384×10^{-5}
κ	0.1	14	0.000267
$-\log(\gamma)$	16.1181	0.10055	0.808
y_0	0.004	52488	4.462×10^{-8}
y_∞	0.002	197531	8.092×10^{-8}
s_0	0.0115	4221	8.45×10^{-7}
s_∞	0.002	84061	3.066×10^{-7}

Table 16: DU: Values for η_0 and V_η prior distributions.

Data Set	θ	κ	$-\log(\gamma)$	y_0	y_∞	s_0	s_∞
DUtt-67223.p.txt 206K 2200/s	0.0589 (0.00711)	0.13 (0.0635)	14.7 (2.58)	0.00519 (0.000404)	0.00248 (0.000956)	0.0121 (0.00143)	0.00529 (0.0016)
DUtt-7510-1.p.txt 198K 0.1/s	0.058 (0.00853)	0.135 (0.0612)	14.9 (2.79)	0.00513 (0.000558)	0.00243 (0.000885)	0.0119 (0.00165)	0.00525 (0.00153)
DUtt023343.p.txt 296K 3400/s	0.0664 (0.0122)	0.136 (0.0657)	14.6 (2.78)	0.00532 (0.000578)	0.0026 (0.00103)	0.0104 (0.00139)	0.00521 (0.00159)
DUtt023283r.p.txt 296K 2800/s	0.0399 (0.0141)	0.117 (0.0654)	15.0 (2.77)	0.00518 (0.000768)	0.00231 (0.000931)	0.0144 (0.00323)	0.00521 (0.00194)
DUtt023100.p.txt 296K 1.0/s	0.0669 (0.0111)	0.14 (0.0637)	14.8 (3.15)	0.00525 (0.000641)	0.0026 (0.000953)	0.0099 (0.00179)	0.00498 (0.00139)
DUtt02310-3b.p.txt 296K 0.001/s	0.0435 (0.00613)	0.129 (0.0558)	15.0 (2.76)	0.00536 (0.000628)	0.00255 (0.000835)	0.0119 (0.00246)	0.00493 (0.00136)
DUtt02310-3br.p.txt 296K 0.001/s	0.0275 (0.00852)	0.119 (0.0543)	15.0 (3.39)	0.0053 (0.000808)	0.00214 (0.000878)	0.0145 (0.00329)	0.00503 (0.00173)
DUtt15010-1.p.txt 423K 0.1/s	0.0181 (0.0024)	0.141 (0.0577)	14.5 (3.71)	0.0057 (0.000875)	0.00248 (0.000842)	0.0125 (0.00389)	0.00385 (0.00109)
DUtt30010-1.p.txt 573K 0.1/s	0.00878 (0.00125)	0.17 (0.0536)	13.6 (3.95)	0.00576 (0.00112)	0.00179 (0.000489)	0.0134 (0.00541)	0.00273 (0.000703)

Table 17: DU: Posterior mean (standard deviation).

and 0.95 quantiles of the MCMC output. The hierarchical model provides a much better fit to the individual data sets than does the common model.

Figure 18 shows 95% posterior probability regions for pairs of parameters. The blue ellipses correspond to the individual data sets. The broken black ellipse and the solid black ellipse correspond to η_0 and η_{K+1} , respectively, from the hierarchical model fit. The cyan ellipse are derived from fitting a common model to all the data. Fitting the common model gives a very optimistic picture of our knowledge about the parameters. Based on these plots, fitting a single nonlinear regression model to all the experimental conditions under-represents the amount of uncertainty present in the PTW parameters. For the plots in the lower triangle the points labeled with a ‘‘p’’ are values given in [1] and the points labeled ‘‘c’’ are given in [2]. This plot was constructed by assuming bivariate normality for the marginal posterior distributions, not accounting for constraints on the individual parameters. Therefore, this plot should be viewed merely as a heuristic for assessing bivariate relationships among the parameter distributions represented.

Figure 19 shows 0.05, 0.50, and 0.95 point-wise posterior prediction intervals based on η_{K+1} and η_0 from the hierarchical model fit. This plot shows how additional posterior uncertainty in η_{K+1} relative to η_0 translates into increased posterior prediction uncertainty.

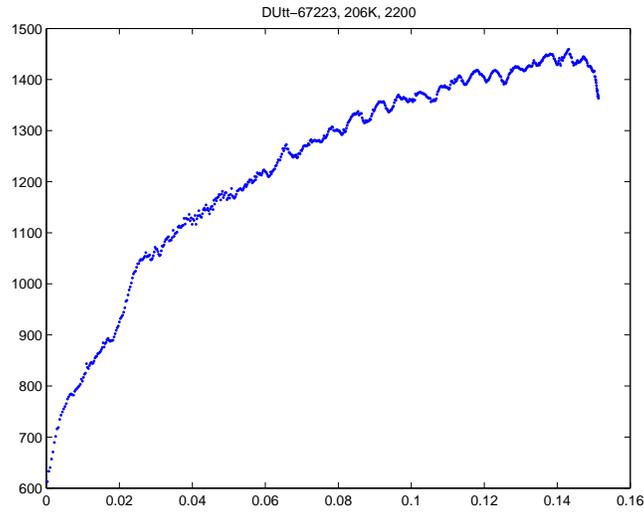


Figure 15: Plot of data obtained from Hopkinson-bar experiment with depleted uranium at 206K and a strain rate of 2200/s.

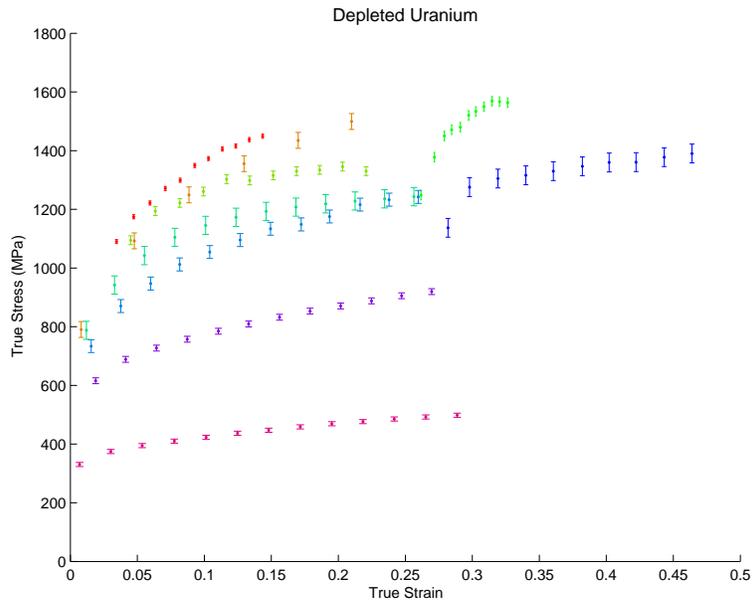


Figure 16: DU: Sampled data with $\pm 1\sigma$ error bars.

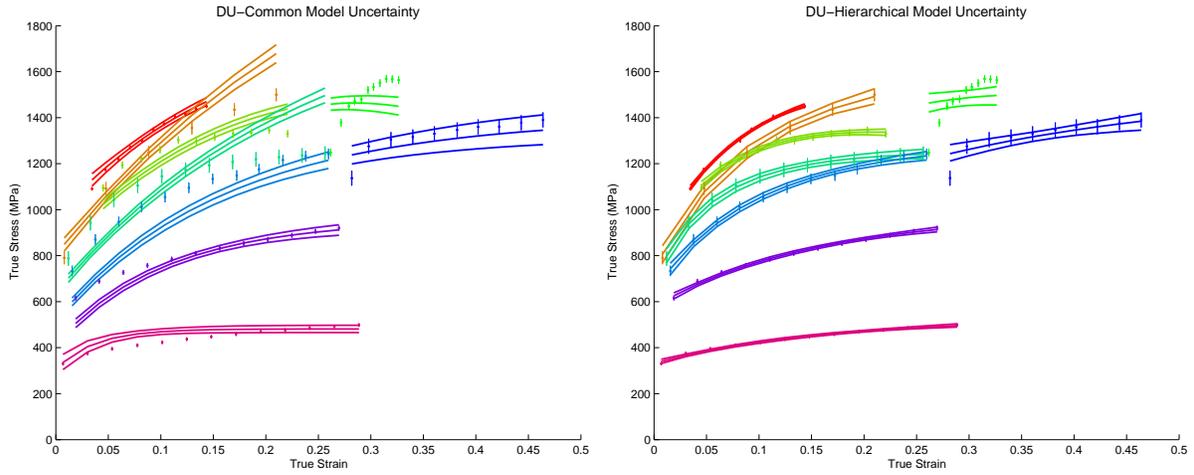


Figure 17: DU: Point-wise posterior prediction intervals from the common model fit, left-hand plot, and the hierarchical model fit, right-hand plot. The curves are 0.05, 0.50, and 0.95 point-wise quantiles from the MCMC output. The curves in the right-hand plot are based on the η_i s.

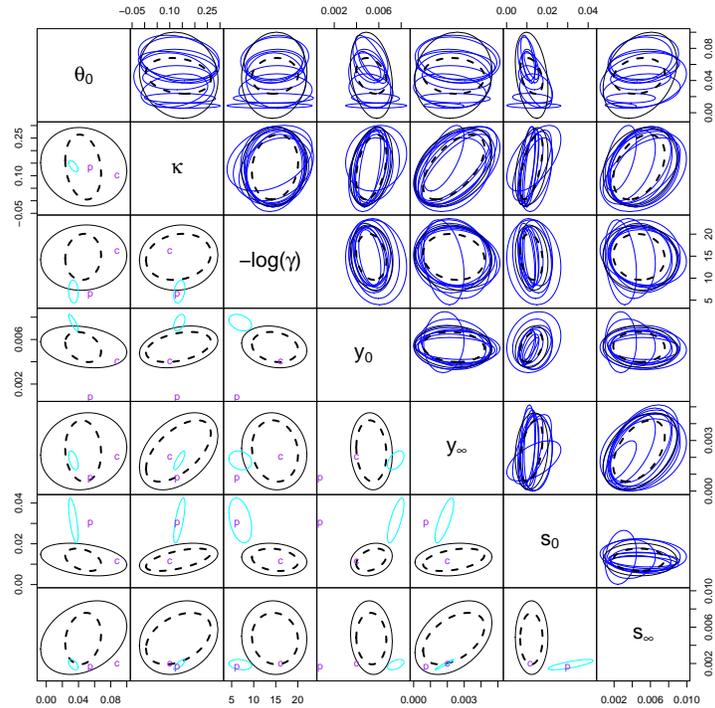


Figure 18: DU: Approximate 95% posterior probability regions. The blue ellipses in the upper right correspond to the η_i s, one for each data set. The broken black ellipses and the solid black ellipses correspond to η_0 and η_{K+1} , respectively, from the hierarchical model fit. The cyan ellipses in the lower left correspond to η_0 from the common model fit.

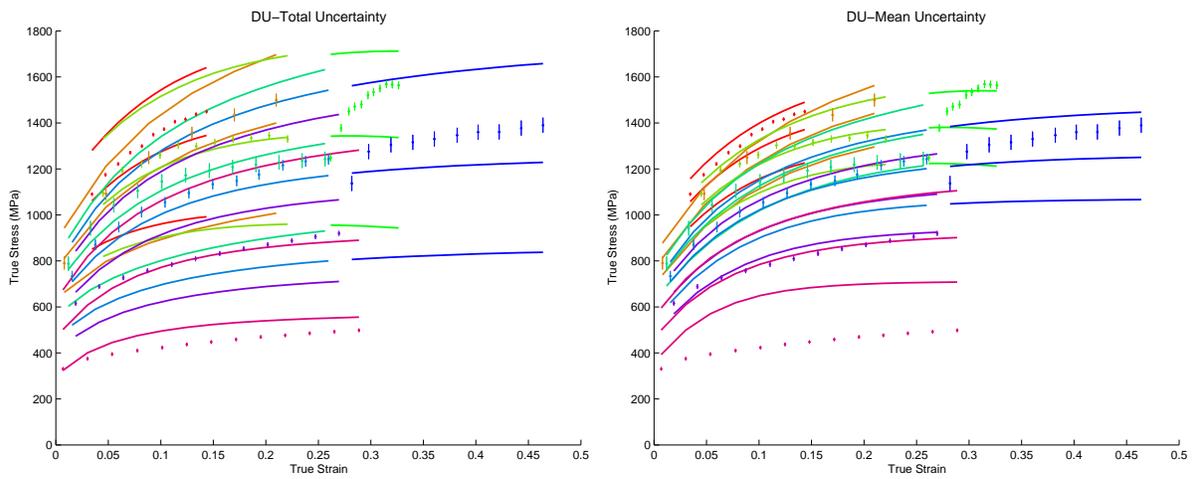


Figure 19: DU: Posterior prediction intervals from the hierarchical model. The left-hand plot is based on $\boldsymbol{\eta}_{K+1}$ and the right-hand plot is based on $\boldsymbol{\eta}_0$. The curves are 0.05, 0.50, and 0.95 point-wise quantiles from the MCMC output.

5.4 Tantalum

In this section we discuss fitting the PTW model to data from Hopkinson-bar experiments on tantalum. There are six tantalum data sets.

Table 18 lists the data sets used along with some of the experimental conditions for each data set. In the table, n_i is the number of observations after thinning and σ_i is the root mean square error (RMSE) from fitting a quadratic model to the thinned data.

Data Set	n_i	Temp (K)	Strain Rate (s^{-1})	σ_i
TaRT133.p.t.txt	29	298	1300	5.8
TaRT10-3.LE.txt	32	298	0.001	10
TaRT10-1.p.txt	10	298	0.1	2
Ta40263.p.t.txt	27	673	2600	6.1
Ta80393.p.t.txt	31	1073	3900	6.3
TaLN10-3.p.txt	13	77	0.001	5

Table 18: Ta: Experimental conditions.

Table 19 gives the values we used for the fixed parameters in the PTW model. These values were provided by [8].

Parameter	Nominal Value	Parameter	Nominal Value
p	4	G_0	722
y_1	0.012	α	0.48
y_2	0.4	C_v	145.5
β	0.23	A	180.95
T_m	3290	ρ	16.6

Table 19: Ta: Fixed PTW parameters.

Table 20 gives the values we used for \mathbf{b}_0 and \mathbf{V}_0 . The values of \mathbf{b}_0 were provided by [8]. This table also gives the prior values we used for $\mathbf{\Omega}$. The appendix describes the method used to determine these values.

Parameter	\mathbf{b}_0	$\mathbf{V}_{0,ii}$	$\mathbf{\Omega}_{ii}$
θ	0.01	0.0	2.583×10^{-6}
κ	0.6	0.0	0.000576
$-\log(\gamma)$	10.127	0.0	0.0941
y_0	0.00925	0.0	5.241×10^{-8}
y_∞	0.00123	0.0	3.843×10^{-9}
s_0	0.0122	0.0	7.263×10^{-6}
s_∞	0.00375	0.0	2.642×10^{-8}

Table 20: Ta: Values for $\boldsymbol{\eta}_0$ and $\mathbf{V}_\boldsymbol{\eta}$ prior distributions.

Table 21 shows the posterior mean and standard deviation for each parameter and each data set. These were calculated from the MCMC output by taking every other realization after discarding the first 100,000 samples. Therefore, these estimates are based on 250,000 samples.

Data Set	θ	κ	$-\log(\gamma)$	y_0	y_∞	s_0	s_∞
TaRT133.p.t.txt 298K 1300/s	0.0122 (0.00225)	0.724 (0.144)	11.5 (1.52)	0.0094 (0.00112)	0.00138 (0.000261)	0.0151 (0.00411)	0.00295 (0.000513)
TaRT10-3.LE.txt 298K 0.001/s	0.00666 (0.00127)	0.695 (0.168)	11.8 (1.76)	0.00929 (0.0014)	0.00139 (0.000289)	0.0185 (0.00578)	0.00347 (0.000585)
TaRT10-1.p.t.txt 298K 0.1/s	0.0127 (0.000965)	0.725 (0.123)	11.4 (1.6)	0.00939 (0.00113)	0.00138 (0.000229)	0.0134 (0.00383)	0.0027 (0.000354)
Ta40263.p.t.txt 673K 2600/s	0.00802 (0.00104)	0.692 (0.164)	11.8 (1.52)	0.00934 (0.00124)	0.00146 (0.000298)	0.0154 (0.00478)	0.00309 (0.000466)
Ta80393.p.t.txt 1073K 3900/s	0.0202 (0.00281)	0.723 (0.148)	11.3 (1.82)	0.00947 (0.00137)	0.0012 (0.000106)	0.0165 (0.00508)	0.00259 (0.000347)
TaLN10-3.p.t.txt 77K 0.001/s	0.0135 (0.00231)	0.71 (0.135)	11.5 (1.63)	0.00949 (0.000837)	0.00135 (0.000234)	0.0137 (0.00158)	0.00288 (0.000504)

Table 21: Ta: Posterior mean (standard deviation).

Figure 20 is a plot of the thinned tantalum data along with $\pm 1\sigma$ error bars.

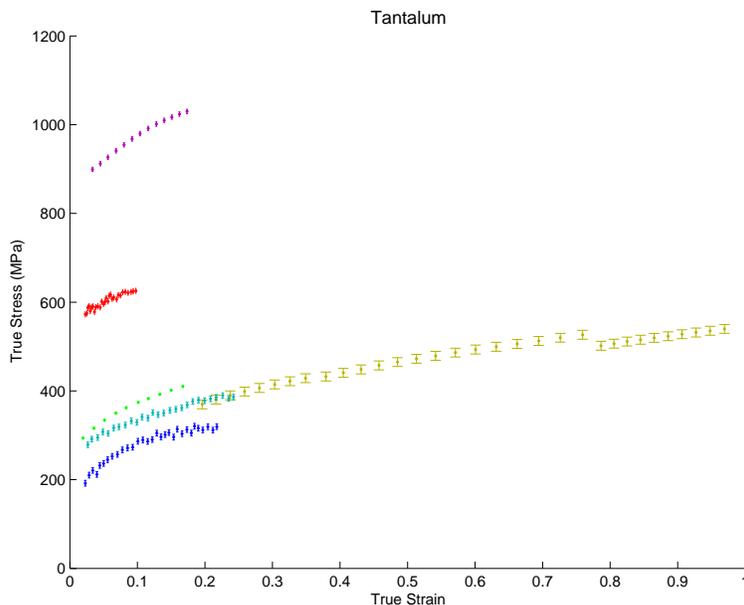


Figure 20: Ta: Sampled data along with $\pm 1\sigma$ error bars.

Figure 21 shows point-wise posterior prediction intervals derived from fitting the common model and the hierarchical model. In both plots these intervals correspond to the 0.05, 0.50, and 0.95 quantiles of the MCMC output. The intervals from the hierarchical model fit are based on the η_i s. The common and hierarchical models provide comparable fits to the individual data sets.

Figure 22 shows approximate 95% posterior probability regions. The blue ellipses correspond to the individual data sets. The broken black ellipses correspond to η_0 and the solid black ellipses correspond to η_{K+1} . The cyan ellipses are derived from fitting a common parameter model to all the data. Fitting the common model gives a somewhat optimistic

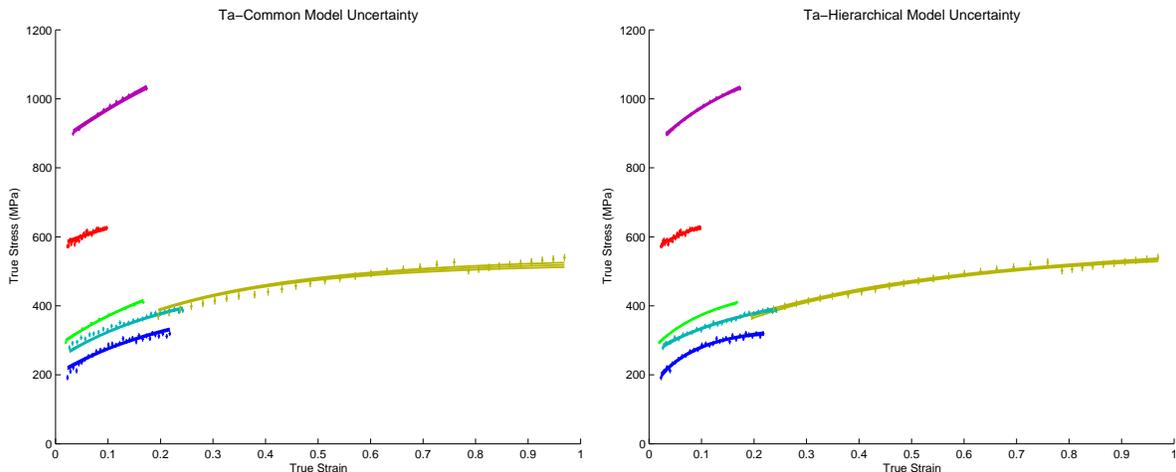


Figure 21: Ta: Point-wise posterior prediction intervals from the common model fit, left-hand plot, and the hierarchical model fit, right-hand plot. The curves are 0.05, 0.50, and 0.95 point-wise quantiles from the MCMC output. The curves in the right-hand plot are based on the η_i s.

picture of our knowledge about the parameters. Based on these plots, fitting a single non-linear regression model to all the experimental conditions under-represents the amount of uncertainty present in the PTW parameters. For the plots in the lower triangle the points labeled with a “p” are point estimates given in [1] and points labeled with a “c” are point estimates given in [2]. Note that all of the ellipse plots were constructed by assuming bivariate normality for the marginal posterior distributions, not accounting for constraints on the individual parameters. Therefore, this plot should be viewed merely as a heuristic for assessing bivariate relationships among the parameter distributions represented.

Also included in this plot are point estimates with a label “h” and approximate 95% posterior probability ellipses (colored green), taken from the analysis of [8]. The model used in [8] is identical to the common model we fit except [8] included an additive systematic bias term for each experimental condition and used non-linear least squares to estimate the PTW parameters and their associated uncertainties. The analysis in [8] used data from three tantalum experiments not incorporated in the analysis of this paper, and this analysis used data from one tantalum experiment not used by [8].

Figure 23 shows 0.05, 0.50, and 0.95 point-wise posterior prediction intervals for η_{K+1} and for η_0 from the hierarchical model fit. This plot shows how additional posterior uncertainty in η_{K+1} relative to η_0 translates into increased posterior prediction uncertainty.

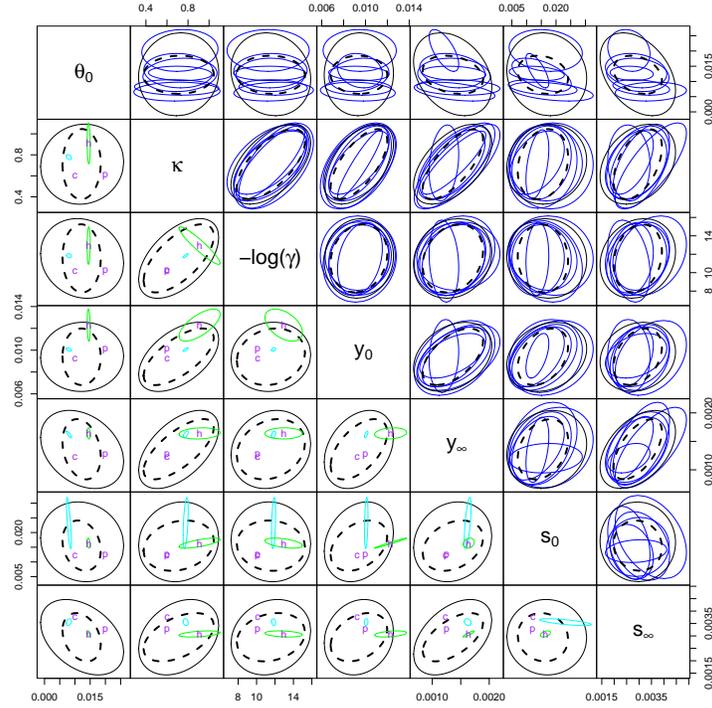


Figure 22: Ta: Approximate 95% posterior probability regions. The blue ellipses in the upper right correspond to the η_i s, one for each data set. The broken black ellipses and the solid black ellipses correspond to η_0 and η_{K+1} , respectively, from the hierarchical model fit. The cyan ellipses in the lower left correspond to η_0 from the common model fit. The green ellipses in the lower left correspond to approximate 95% posterior probability regions from fitting the common model with an additive systematic bias term for each experimental condition, as described in [8].

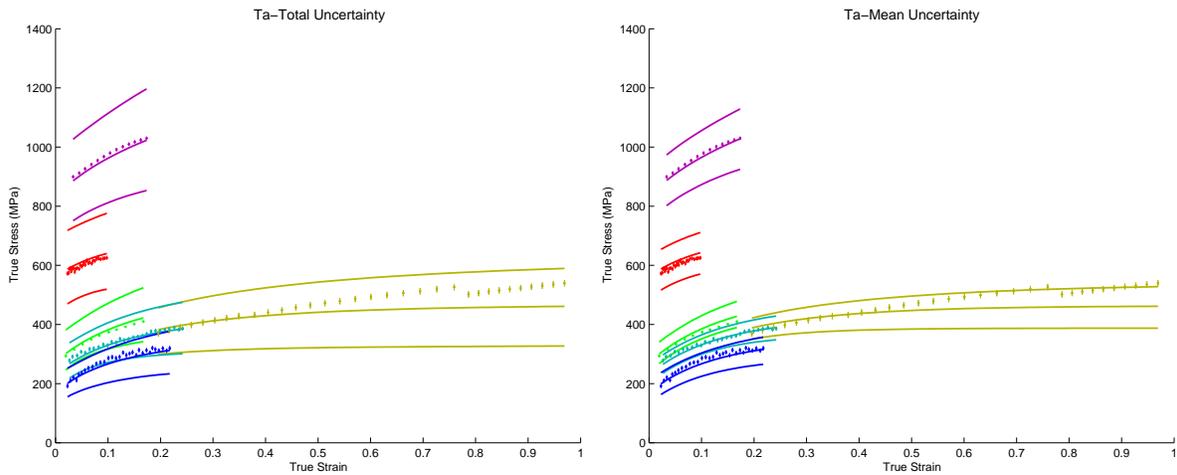


Figure 23: Ta: Posterior prediction intervals from the hierarchical model. The left-hand plot is based on η_{K+1} and the right-hand plot is based on η_0 . The curves are 0.05, 0.50, and 0.95 point-wise quantiles from the MCMC output.

5.5 Uranium 6 wt% Niobium

In this section we discuss fitting the PTW model to data from Hopkinson-bar experiments on uranium 6 wt% niobium. There are ten U-6Nb data sets.

Table 22 lists the data sets used along with some of the experimental conditions for each data set. In the table, n_i is the number of observations after thinning and σ_i is the root mean square error (RMSE) from fitting a quadratic model to the thinned data.

Data Set	n_i	Temp (K)	Strain Rate (s^{-1})	σ_i
U6LN10-3.p.txt	8	77	0.001	81.0
U6b.173.1300.txt	29	173	1300	58.0
U6.173.1e-3.LE.p.txt	21	173	0.001	57.0
U6b.233.1e0.LE.p.txt	14	233	1.0	57.0
U6b.298.1400.txt	29	298	1400	58.0
U620253.txt	19	473	2500	40.0
U6rt10-3c.txt	13	298	0.001	43.0
U630303.txt	12	573	3000	45.0
U640303.txt	15	673	3000	13.0
U660323.txt	14	873	3200	11.0

Table 22: U-6Nb: Experimental conditions.

Table 23 gives the values we used for the fixed parameters in the PTW model. These values were provided by [2].

Parameter	Nominal Value	Parameter	Nominal Value
p	5.0	G_0	1018
y_1	0.0515	α	0.8057
y_2	0.27	C_v	117.5
β	0.27	A	218.4542
T_m	1475	ρ	17.38

Table 23: U-6Nb: Fixed PTW parameters.

Table 24 gives the values we used for \mathbf{b}_0 and \mathbf{V}_0 . The values of \mathbf{b}_0 were provided by [2]. This table also gives the values we used for $\mathbf{\Omega}$. The appendix describes the method used to determine these values.

Table 25 shows the posterior mean and standard deviation for each parameter and each data set. These were calculated from the MCMC output by taking every other realization after discarding the first 100,000 samples. Therefore, these estimates are based on 250,000 samples.

Figure 24 is a plot of the thinned data used for fitting along with $\pm 1\sigma$ error bars.

Figure 25 shows point-wise posterior prediction intervals derived from fitting the common model and the hierarchical model. In both plots these intervals correspond to the 0.05, 0.50, and 0.95 quantiles of the MCMC output. The hierarchical model provides a much better fit to the individual data sets than does the common model.

Parameter	b_0	$V_{0,ii}$	Ω_{ii}
θ	0.08	1069	0.000149
κ	0.40	14	0.00157
$-\log(\gamma)$	5.809	0.10055	1.544
y_0	0.00519	52488	2.204×10^{-6}
y_∞	0.00186	197531	1.105×10^{-7}
s_0	0.0515	4221	5.515×10^{-5}
s_∞	0.0105	84061	3.578×10^{-8}

Table 24: U-6Nb: Values for η_0 and V_η prior distributions.

Data Set	θ	κ	$-\log(\gamma)$	y_0	y_∞	s_0	s_∞
U6LN10-3.p.txt 77K 0.001/s	0.166 (0.0528)	0.61 (0.125)	1.87 (4.73)	0.00619 (0.00339)	0.000655 (0.000485)	0.0749 (0.0328)	0.00614 (0.000553)
U6b.173.1300.txt 173K 1300/s	0.136 (0.0118)	0.525 (0.077)	3.9 (2.67)	0.00391 (0.00218)	0.000308 (0.000234)	0.0838 (0.0197)	0.00633 (0.000346)
U6.173.1e-3.LE.p.txt 173K 0.001/s	0.0744 (0.00803)	0.447 (0.0488)	8.27 (2.48)	0.00682 (0.00318)	0.00071 (0.000283)	0.0705 (0.0171)	0.00662 (0.000406)
U6b.233.1e0.LE.p.txt 233K 1.0/s	0.119 (0.013)	0.53 (0.0781)	4.71 (2.69)	0.0058 (0.00284)	0.000346 (0.000249)	0.0776 (0.0192)	0.0064 (0.000275)
U6b.298.1400.txt 298K 1400/s	0.165 (0.0214)	0.53 (0.121)	2.58 (3.63)	0.00507 (0.00303)	0.000384 (0.000276)	0.0865 (0.0258)	0.00615 (0.000295)
U620253.txt 473K 2500/s	0.113 (0.0156)	0.514 (0.0915)	4.15 (2.52)	0.00598 (0.00292)	0.000453 (0.000282)	0.0754 (0.0201)	0.00621 (0.000153)
U6rt10-3c.txt 298K 0.001/s	0.0844 (0.00924)	0.49 (0.0876)	7.02 (2.61)	0.00652 (0.00316)	0.00033 (0.000239)	0.0725 (0.0175)	0.00659 (0.000232)
U630303.txt 573K 3000/s	0.127 (0.0376)	0.526 (0.104)	4.78 (4.06)	0.00612 (0.00319)	0.000759 (0.000462)	0.073 (0.0231)	0.00605 (0.000239)
U640303.txt 673K 3000/s	0.0177 (0.0022)	0.453 (0.11)	11.2 (1.96)	0.00821 (0.00449)	0.000984 (0.000743)	0.0608 (0.0229)	0.00684 (0.000513)
U660323.txt 873K 3200/s	0.0313 (0.00725)	0.45 (0.113)	10.6 (1.83)	0.00792 (0.00424)	0.000913 (0.000685)	0.0633 (0.0211)	0.0069 (0.000456)

Table 25: U-6Nb: Posterior mean (standard deviation).

Figure 26 shows approximate 95% posterior probability regions. The blue ellipses correspond to the individual data sets. The broken black ellipses and the solid black ellipses correspond to η_0 and η_{K+1} , respectively, from the hierarchical model fit. The cyan ellipses are derived from fitting the common model to all the data. Fitting the common model gives a very optimistic picture of our knowledge about the parameters. Based on these plots, fitting a single nonlinear regression model to all the experimental conditions under-represents the amount of uncertainty present in the PTW parameters. For the plots in the lower triangle the three points labeled with “c” represent three distinct sets of point estimates given in [2]. This plot was constructed by assuming bivariate normality for the marginal posterior distributions, not accounting for constraints on the individual parameters. Therefore, this plot should be viewed merely as a heuristic for assessing bivariate relationships among the parameter distributions represented.

Figure 27 shows 0.05, 0.50, and 0.95 point-wise posterior prediction intervals for η_{K+1} and for η_0 from the hierarchical model fit. This plot shows how additional posterior uncertainty in η_{K+1} relative to η_0 translates into increased posterior prediction uncertainty.

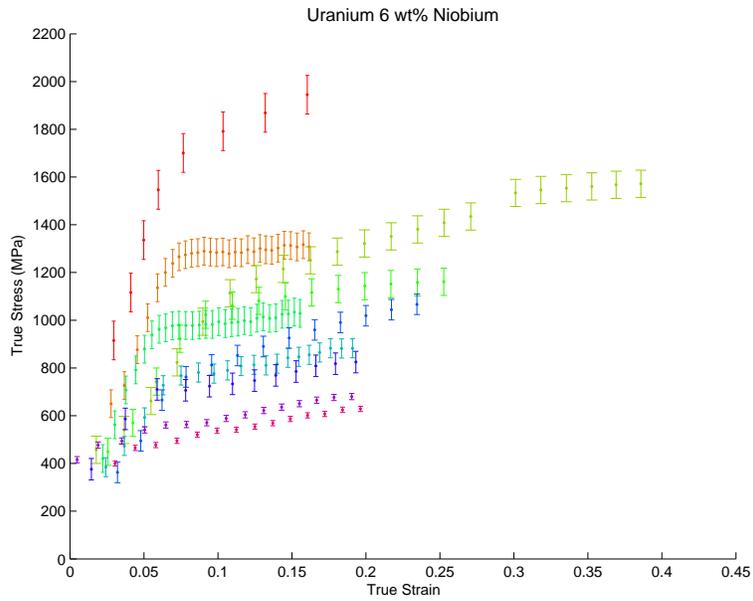


Figure 24: U-6Nb: Sampled data along with $\pm 1\sigma$ error bars.

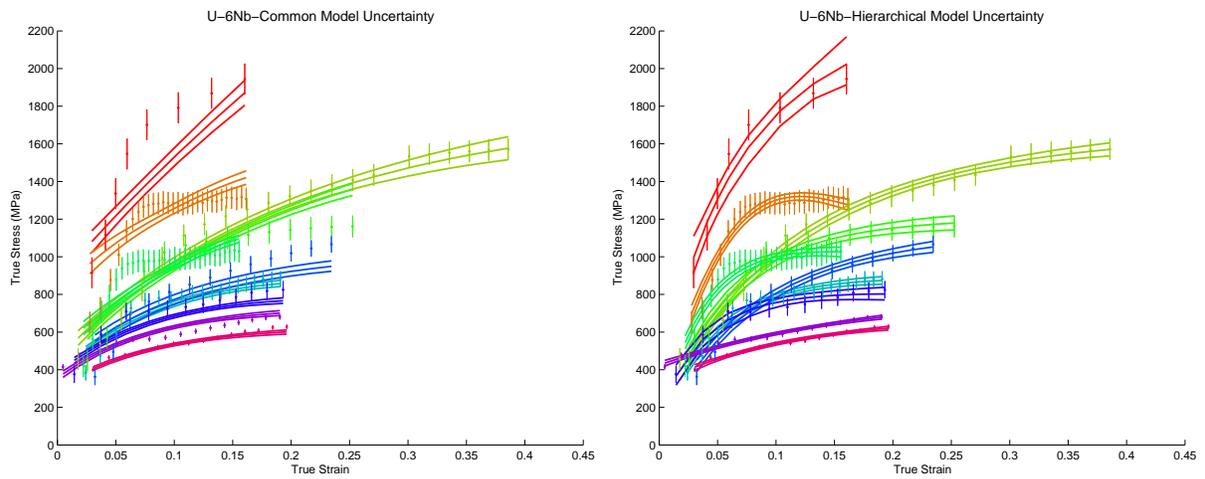


Figure 25: U-6Nb: Point-wise posterior prediction intervals from the common model fit, left-hand plot, and the hierarchical model fit, right-hand plot. The curves are 0.05, 0.50, and 0.95 point-wise quantiles from the MCMC output. The curves in the right-hand plot are based on the $\boldsymbol{\eta}_i$ s.

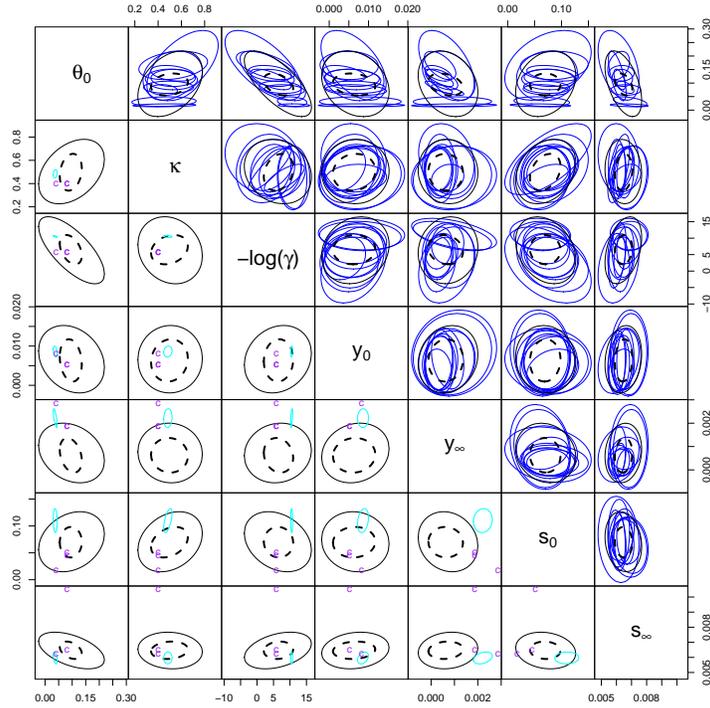


Figure 26: U-6Nb: Approximate 95% posterior probability regions. The blue ellipses in the upper right correspond to the η_i s, one for each data set. The broken black ellipses and the solid black ellipses correspond to η_0 and η_{K+1} , respectively, from the hierarchical model fit. The cyan ellipses in the lower left correspond to η_0 from the common model fit.

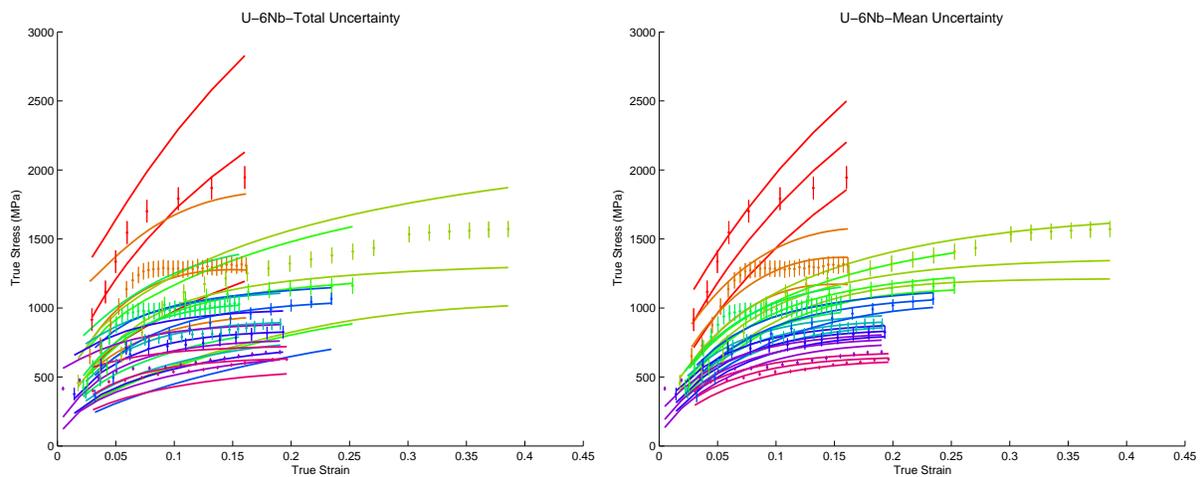


Figure 27: U-6Nb: Posterior prediction intervals from the hierarchical model. The left-hand plot is based on η_{K+1} and the right-hand plot is based on η_0 . The curves are 0.05, 0.50, and 0.95 point-wise quantiles from the MCMC output.

6 Discussion

The uncertainty analysis for the PTW model parameters presented in this paper is directed at obtaining a realistic quantification of the uncertainty in these parameters for several materials across a range of temperatures and strain rates representative of data currently available from Hopkinson-bar and quasi-static-compression experiments. The relevant uncertainty is summarized by the posterior predictive distribution $p(\boldsymbol{\eta}_{K+1} | \mathbf{y})$ of equation (6). Section 3 explains the relevance of this uncertainty quantification to dynamic experiments in which strain rates and temperatures vary over a potentially wide range of values. However, this distribution will likely over-represent the uncertainty in the PTW model parameters, if prediction is to be specialized to a particular temperature and strain rate environment. One possible revision to equation (6) for uncertainty quantification in this setting involves specifying a set of weights $\{w_i\}$, $i = 1, \dots, N$, for which $\sum_{i=1}^N w_i = 1$, and computing

$$p(\boldsymbol{\eta}_{K+1}^w | \mathbf{y}) \sim \sum_{i=1}^N w_i p(\boldsymbol{\eta}_{K+1} | \boldsymbol{\eta}_{0,i}, \mathbf{V}_{\boldsymbol{\eta},i}).$$

Methodology for determining the weights is required. As a general principle, the relative ordering of the weights would be determined by the degree to which the corresponding samples of $\boldsymbol{\eta}_0$ are consistent with the temperature and strain rate at which prediction is desired.

These uncertainties should serve as a starting point for using data from Taylor cylinder or flyer plate tests to further constrain parameter uncertainties. Such an analysis will likely have to be coupled with detailed simulation code to obtain the resulting inference. See [5] for such an example.

A Appendix

In this appendix we discuss more fully the MCMC algorithm used to generate samples from the posterior distribution of the parameters, assuming the model and prior distributions given in (4). This posterior distribution is not analytically tractable. Therefore, we use the Gibbs sampler ([10]) to generate samples from this posterior distribution. The Gibbs sampler requires the full conditional distribution for each parameter. The full conditionals that we need are given in (5).

The full conditional posterior distribution of $\boldsymbol{\eta}$ is proportional to the data likelihood times a prior and is written as $p(\boldsymbol{\eta} | \mathbf{y}, \boldsymbol{\xi}, \boldsymbol{\eta}_0, \mathbf{V}_\boldsymbol{\eta}) \propto p(\mathbf{y} | \boldsymbol{\eta}, \boldsymbol{\xi}) p(\boldsymbol{\eta} | \boldsymbol{\eta}_0, \mathbf{V}_\boldsymbol{\eta})$. For the prior distributions given in equation (4) we have:

$$p(\boldsymbol{\eta} | \mathbf{y}, \boldsymbol{\xi}, \boldsymbol{\eta}_0, \mathbf{V}_\boldsymbol{\eta}) \propto \prod_{i=1}^K \exp \left\{ -\frac{1}{2} [(\mathbf{y}_i - \mathbf{g}_i(\boldsymbol{\eta}_i))^\top \boldsymbol{\Sigma}_i(\boldsymbol{\xi}_i)^{-1} (\mathbf{y}_i - \mathbf{g}_i(\boldsymbol{\eta}_i)) + (\boldsymbol{\eta}_i - \boldsymbol{\eta}_0)^\top \mathbf{V}_\boldsymbol{\eta} (\boldsymbol{\eta}_i - \boldsymbol{\eta}_0)] \right\} I(\boldsymbol{\eta}_i \in C).$$

To sample this posterior distribution we use a Metropolis step ([11], [12]).

The covariance matrix of the data depends on the vector $\boldsymbol{\xi}$, which is typically unknown. The full conditional posterior distribution of $\boldsymbol{\xi}$ is written as $p(\boldsymbol{\xi} | \mathbf{y}, \boldsymbol{\eta}, \boldsymbol{\eta}_0, \mathbf{V}_\boldsymbol{\eta}) \propto p(\mathbf{y} | \boldsymbol{\eta}, \boldsymbol{\xi}) p(\boldsymbol{\xi})$, which is again the data likelihood times a prior. Using the distributions given in equation (4) and denoting the prior distribution of $\boldsymbol{\xi}_i$ as $p(\boldsymbol{\xi}_i)$, the full conditional posterior is

$$p(\boldsymbol{\xi} | \mathbf{y}, \boldsymbol{\eta}, \boldsymbol{\eta}_0, \mathbf{V}_\boldsymbol{\eta}) \propto \prod_{i=1}^K \frac{p(\boldsymbol{\xi}_i)}{\sqrt{\det[\boldsymbol{\Sigma}_i(\boldsymbol{\xi}_i)]}} \exp \left\{ -\frac{1}{2} (\mathbf{y}_i - \mathbf{g}_i(\boldsymbol{\eta}_i))^\top \boldsymbol{\Sigma}_i(\boldsymbol{\xi}_i)^{-1} (\mathbf{y}_i - \mathbf{g}_i(\boldsymbol{\eta}_i)) \right\},$$

where $\det[\cdot]$ denotes the determinant of the matrix argument. The results presented in subsequent sections assume the measurement error model (1), corresponding to setting $\boldsymbol{\xi}_i = \lambda_i$ and $\boldsymbol{\Sigma}_i(\boldsymbol{\xi}_i) = \frac{1}{\lambda_i} \text{Diag}(\sigma_i^2; n_i)$. The notation $\text{Diag}(\sigma_i^2; n_i)$ indicates a $n_i \times n_i$ diagonal matrix with σ_i^2 on the diagonal.

Assuming a $\text{Gamma}(a_i, b_i)$ prior distribution for λ_i , we have $p(\lambda_i) \propto \lambda_i^{a_i-1} \exp\{-b_i \lambda_i\}$. The full conditional posterior distribution of λ_i is proportional to the data likelihood times this prior and we write $p(\lambda_i | \mathbf{y}_i, \boldsymbol{\eta}_i, \boldsymbol{\eta}_0, \mathbf{V}_\boldsymbol{\eta}) \propto p(\mathbf{y}_i | \boldsymbol{\eta}_i, \lambda_i) p(\lambda_i)$. Collecting terms we have

$$p(\lambda_i | \mathbf{y}_i, \boldsymbol{\eta}_i, \boldsymbol{\eta}_0, \mathbf{V}_\boldsymbol{\eta}) \sim \text{Gamma} \left(a_i + \frac{n_i}{2}, b_i + \frac{1}{2} (\mathbf{y}_i - \mathbf{g}_i(\boldsymbol{\eta}_i))^\top \text{Diag}(\sigma_i^2; n_i)^{-1} (\mathbf{y}_i - \mathbf{g}_i(\boldsymbol{\eta}_i)) \right).$$

A Gibbs update is used to sample λ_i from this conditional posterior distribution.

The fixed quantity σ_i^2 is estimated by fitting a quadratic model to the (sub-sampled) data from experiment i and then computing the mean square error (MSE) of the residuals. This MSE is then used for σ_i^2 . The parameter λ_i controls the amount of allowable deviation from the estimated data error σ_i^2 . The $\text{Gamma}(a_i, b_i)$ prior distribution for λ_i will generally be

chosen to satisfy $a_i = b_i$, having mean 1 and variance $1/a_i$, so that increasing a_i results in a smaller allowable deviation from σ_i^2 .

The full conditional posterior distribution of $\boldsymbol{\eta}_0$ is again proportional to a likelihood times a prior but now the likelihood is $p(\boldsymbol{\eta} | \boldsymbol{\eta}_0, \mathbf{V}_\boldsymbol{\eta})$. Combining this likelihood with the prior for $\boldsymbol{\eta}_0$ gives $p(\boldsymbol{\eta}_0 | \mathbf{y}, \boldsymbol{\eta}, \boldsymbol{\xi}, \mathbf{V}_\boldsymbol{\eta}) \propto p(\boldsymbol{\eta} | \boldsymbol{\eta}_0, \mathbf{V}_\boldsymbol{\eta}) p(\boldsymbol{\eta}_0)$. Collecting terms we have

$$p(\boldsymbol{\eta}_0 | \mathbf{y}, \boldsymbol{\eta}, \boldsymbol{\xi}, \mathbf{V}_\boldsymbol{\eta}) \sim \mathcal{N}((K\mathbf{V}_\boldsymbol{\eta} + \mathbf{V}_0)^{-1}(K\mathbf{V}_\boldsymbol{\eta}\bar{\boldsymbol{\eta}} + \mathbf{V}_0\mathbf{b}_0), (K\mathbf{V}_\boldsymbol{\eta} + \mathbf{V}_0)^{-1}) I(\boldsymbol{\eta}_0 \in C). \quad (8)$$

Here $\bar{\boldsymbol{\eta}} = \frac{1}{K} \sum_{i=1}^K \boldsymbol{\eta}_i$. We use a Gibbs update to sample this conditional posterior distribution.

The full conditional posterior distribution of $\mathbf{V}_\boldsymbol{\eta}$ is proportional to the $\boldsymbol{\eta}$ likelihood times the prior for $\mathbf{V}_\boldsymbol{\eta}$ and we write $p(\mathbf{V}_\boldsymbol{\eta} | \mathbf{y}, \boldsymbol{\eta}, \boldsymbol{\xi}, \boldsymbol{\eta}_0) \propto p(\boldsymbol{\eta} | \boldsymbol{\eta}_0, \mathbf{V}_\boldsymbol{\eta}) p(\mathbf{V}_\boldsymbol{\eta})$. Combining terms we have

$$p(\mathbf{V}_\boldsymbol{\eta} | \mathbf{y}, \boldsymbol{\eta}, \boldsymbol{\xi}, \boldsymbol{\eta}_0) \sim W \left[\left(\sum_{i=1}^K (\boldsymbol{\eta}_i - \boldsymbol{\eta}_0)(\boldsymbol{\eta}_i - \boldsymbol{\eta}_0)^\top + \nu\boldsymbol{\Omega} \right)^{-1}, \nu + K \right].$$

A Gibbs update is used to sample this conditional posterior distribution.

The MCMC algorithm is completed with a decorrelating Metropolis step, implemented as follows:

1. set $j = 1$
2. generate z_j from a symmetric distribution
3. propose the candidate $\{\boldsymbol{\eta}_0 + z_j\mathbf{e}_j, \boldsymbol{\eta}_1 + z_j\mathbf{e}_j, \dots, \boldsymbol{\eta}_K + z_j\mathbf{e}_j\}$, where \mathbf{e}_j denotes the unit vector with 1 in the j th element
4. evaluate the candidate with respect to the joint conditional posterior

$$\begin{aligned} p(\boldsymbol{\eta}_0, \boldsymbol{\eta} | \mathbf{y}, \boldsymbol{\xi}, \mathbf{V}_\boldsymbol{\eta}) &\propto \\ &\prod_{i=1}^K \exp \left\{ -\frac{1}{2}(\mathbf{y}_i - \mathbf{g}_i(\boldsymbol{\eta}_i))^\top \boldsymbol{\Sigma}_i(\boldsymbol{\xi}_i)^{-1}(\mathbf{y}_i - \mathbf{g}_i(\boldsymbol{\eta}_i)) \right\} \times \\ &\prod_{i=1}^K \exp \left\{ -\frac{1}{2}(\boldsymbol{\eta}_i - \boldsymbol{\eta}_0)^\top \mathbf{V}_\boldsymbol{\eta}(\boldsymbol{\eta}_i - \boldsymbol{\eta}_0) \right\} I(\boldsymbol{\eta}_i \in C) \times \\ &\exp \left\{ -\frac{1}{2}(\boldsymbol{\eta}_0 - \mathbf{b}_0)^\top \mathbf{V}_0(\boldsymbol{\eta}_0 - \mathbf{b}_0) \right\} I(\boldsymbol{\eta}_0 \in C) \end{aligned}$$

5. scan through $j = 2, \dots, |\boldsymbol{\eta}_0|$, where $|\boldsymbol{\eta}_0|$ denotes the number of elements in $\boldsymbol{\eta}_0$

This step improves the mixing of components in $\boldsymbol{\eta}_0$ and $\boldsymbol{\eta}_i$, $i = 1, \dots, K$, that are subject to a high degree of shrinkage (i.e., components of $\boldsymbol{\eta}_i$ that are nearly identical to the corresponding components of $\boldsymbol{\eta}_0$ in a probabilistic sense).

Inferences about a new parameter vector $\boldsymbol{\eta}_{K+1}$ are drawn from the posterior distribution $p(\boldsymbol{\eta}_{K+1} | \mathbf{y})$, which is

$$\begin{aligned} p(\boldsymbol{\eta}_{K+1} | \mathbf{y}) &= \int p(\boldsymbol{\eta}_{K+1} | \boldsymbol{\eta}_0, \mathbf{V}_\boldsymbol{\eta}) p(\boldsymbol{\eta}_0, \mathbf{V}_\boldsymbol{\eta} | \mathbf{y}) d\boldsymbol{\eta}_0 d\mathbf{V}_\boldsymbol{\eta} \\ &\approx \frac{1}{N} \sum_{j=1}^N p(\boldsymbol{\eta}_{K+1} | \boldsymbol{\eta}_{0,j}, \mathbf{V}_{\boldsymbol{\eta},j}). \end{aligned} \quad (9)$$

The unconstrained version of distribution (9) has mean vector $\bar{\boldsymbol{\eta}}_0 = \frac{1}{N} \sum_{j=1}^N \boldsymbol{\eta}_{0,j}$ and covariance matrix

$$\frac{1}{N} \sum_{j=1}^N [\mathbf{V}_{\boldsymbol{\eta},j}^{-1} + (\boldsymbol{\eta}_{0,j} - \bar{\boldsymbol{\eta}}_0)(\boldsymbol{\eta}_{0,j} - \bar{\boldsymbol{\eta}}_0)^\top].$$

The summation is over the MCMC samples. The structure of this covariance matrix indicates two sources of uncertainty: variability in parameter vector samples across individual data sets represented by $\mathbf{V}_\boldsymbol{\eta}^{-1}$, and variability in estimating the population mean $\boldsymbol{\eta}_0$. This posterior distribution gives us information about values of $\boldsymbol{\eta}$ that are representative of the population uncertainty in the PTW parameters across the range of strain rates and temperatures observed in the experimental data used to train the hierarchical model.

We can also draw inference about the population mean $\boldsymbol{\eta}_0$ through its posterior distribution $p(\boldsymbol{\eta}_0 | \mathbf{y})$,

$$p(\boldsymbol{\eta}_0 | \mathbf{y}) \approx \frac{1}{N} \sum_{j=1}^N p(\boldsymbol{\eta}_0 | \mathbf{y}, \boldsymbol{\eta}_j, \mathbf{V}_{\boldsymbol{\eta},j}). \quad (10)$$

Applying the assumptions of (4), it can be shown that

$$p(\boldsymbol{\eta}_0 | \mathbf{y}, \boldsymbol{\eta}, \mathbf{V}_\boldsymbol{\eta}) = \frac{p(\boldsymbol{\eta} | \boldsymbol{\eta}_0, \mathbf{V}_\boldsymbol{\eta}) p(\boldsymbol{\eta}_0) p(\mathbf{V}_\boldsymbol{\eta})}{p(\boldsymbol{\eta}, \mathbf{V}_\boldsymbol{\eta})} = p(\boldsymbol{\eta}_0 | \mathbf{y}, \boldsymbol{\eta}, \boldsymbol{\xi}, \mathbf{V}_\boldsymbol{\eta}).$$

The full conditional distribution $p(\boldsymbol{\eta}_0 | \mathbf{y}, \boldsymbol{\eta}, \boldsymbol{\xi}, \mathbf{V}_\boldsymbol{\eta})$ is given in (8). Let $\tilde{\boldsymbol{\eta}} = (K\mathbf{V}_\boldsymbol{\eta} + \mathbf{V}_0)^{-1}(K\mathbf{V}_\boldsymbol{\eta}\bar{\boldsymbol{\eta}} + \mathbf{V}_0\mathbf{b}_0)$. The unconstrained version of the distribution (10) has mean vector $\bar{\tilde{\boldsymbol{\eta}}} = \frac{1}{N} \sum_{j=1}^N \tilde{\boldsymbol{\eta}}_j$ and covariance matrix

$$\frac{1}{N} \sum_{j=1}^N \left[(K\mathbf{V}_{\boldsymbol{\eta},j} + \mathbf{V}_0)^{-1} + (\tilde{\boldsymbol{\eta}}_j - \bar{\tilde{\boldsymbol{\eta}}})(\tilde{\boldsymbol{\eta}}_j - \bar{\tilde{\boldsymbol{\eta}}})^\top \right].$$

Posterior uncertainty in $\boldsymbol{\eta}_0$ arises from two sources: between sample and prior variability represented by $(K\mathbf{V}_\boldsymbol{\eta} + \mathbf{V}_0)^{-1}$, and variability in estimating the mean of $\boldsymbol{\eta}_0$.

Initial values for the $\boldsymbol{\eta}_i$ s are required to start the MCMC. For a given material, these values were obtained by adding a small amount of white noise to the nominal values \mathbf{b}_0 for that material. We chose the prior distribution of $\boldsymbol{\eta}_0$ to be normal with mean \mathbf{b}_0 and precision matrix \mathbf{V}_0 , subject to the constraint $\boldsymbol{\eta}_0 \in C$. The values of \mathbf{b}_0 for each material were taken from [2]. For Ta and Be, \mathbf{V}_0 was taken to be the zero matrix. For DU, Al and U-6Nb, \mathbf{V}_0 was a diagonal matrix with the sample precisions calculated from the values of \mathbf{b}_0 for all five materials.

We chose the prior distribution for the precision matrix $\mathbf{V}_\boldsymbol{\eta}$ to be Wishart with parameters ν and $\boldsymbol{\Omega}$. For every material but Al, the parameter ν was set to the smallest possible value

ensuring a proper prior distribution, the number of free PTW parameters, which is 7. For Al, ν took the value 10 to strengthen prior information about \mathbf{V}_η as only four data sets were available for analysis. For Al, DU, and U-6Nb, the $\mathbf{\Omega}$ matrix was obtained using the following procedure: Let $\mathbf{\Sigma}_{KH}$ denote the estimated covariance matrix from the common model fit of the PTW model ([8]) to Beryllium experiments. Let \mathbf{D}_1 be the diagonal matrix with elements equal to the diagonal elements of $\mathbf{\Sigma}_{KH}$. Run a pilot MCMC on the complete hierarchical model for 10,000 iterations with $\mathbf{\Omega}$ taken to be \mathbf{D}_1 . From this pilot run we compute the sample covariance matrix of the 10,000 η_0 s and call it \mathbf{S}_1 . We then set $\mathbf{\Omega}$ to be

$$\mathbf{\Omega} = 0.5 * \text{Diag}(\mathbf{S}_1)$$

where $\text{Diag}(\mathbf{S}_1)$ is the diagonal matrix consisting of the diagonal elements of \mathbf{S}_1 . The matrix \mathbf{S}_1 is material dependent. For Ta and Be, we use the same procedure where the material specific $\mathbf{\Sigma}_{KH}$ is obtained from the common model fit of [8]. The Wishart distribution of equation (4) has $E(\mathbf{V}_\eta) = \mathbf{\Omega}^{-1}$ and $\text{Var}(\mathbf{V}_{\eta,ij}) = ((\mathbf{\Omega}^{ij})^2 + \mathbf{\Omega}^{ii}\mathbf{\Omega}^{jj})$, where $\mathbf{\Omega}^{ij}$ denotes the (i, j) entry of $\mathbf{\Omega}^{-1}$. Therefore, the coefficient 0.5 in the definition of $\mathbf{\Omega}$ imposes additional prior shrinkage of the η_i s and inflates the prior variance of each element in \mathbf{V}_η by a factor of 4 relative to using a coefficient of unity, which builds in some vagueness to the prior specification of \mathbf{V}_η .

Acknowledgments

We acknowledge Ralph Nelson for his insights and leadership of the uncertainty quantification project that generated this extension of existing methodology for assessing uncertainty in material model parameters used in dynamic simulations.

References

- [1] Dean L. Preston, Davis L. Tonks, and Duane C. Wallace. Model of plastic deformation for extreme loading conditions. *J. Appl. Phys.*, 93:211–220, 2003.
- [2] Shuh-Rong Chen and G. T. Gray III. Summary of the PTW model parameters. *Los Alamos National Laboratory Report*, LA-CP-04-0920, 2004.
- [3] K. M. Hanson. Introduction to PTW. *Los Alamos National Laboratory Report*, LA-UR-04-0305, 2004.
- [4] Duane C. Wallace. Irreversible thermodynamics of overdriven shocks in solids. *Phys. Rev. B*, 24:5597–5606, 1981.
- [5] Brian Williams, Dave Higdon, Jim Gattiker, Leslie Moore, Michael McKay, and Sallie Keller-McNulty. Combining experimental data and computer simulations, with an application to flyer plate experiments. *Bayesian Analysis*, to appear, 2006.
- [6] Andrew Gelman, John B. Carlin, Hal S. Stern, and Donald B. Rubin. *Bayesian Data Analysis*. Chapman & Hall, Boca Raton, Florida, 1995.
- [7] Bradley P. Carlin and Thomas A. Louis. *Bayes and Empirical Bayes Methods for Data Analysis*. Chapman & Hall, Boca Raton, Florida, 2000.
- [8] K. M. Hanson. Inference about the plastic behavior of materials from experimental data. In K. M. Hanson and F. M. Hemez, editors, *Sensitivity Analysis of Model Output*, pages 126–136. Los Alamos Research Library, Los Alamos, New Mexico, 2005.
- [9] José C. Pinheiro and Douglas M. Bates. *Mixed-effects Models in S and S-PLUS*. Springer, New York, 2000.
- [10] George Casella and Edward I. George. Explaining the Gibbs sampler. *The American Statistician*, 46, 1992.
- [11] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equations of state calculations by fast computing machine. *J. Chem. Phys.*, 21:1087–1091, 1953.
- [12] Siddhartha Chib and Edward Greenberg. Understanding the Metropolis–Hastings algorithm. *The American Statistician*, 49, 1995.