SCALE-BRIDGING MODEL DEVELOPMENT AND INCREASED MODEL CREDIBILITY

by

Benjamin Brian Schroeder

A dissertation submitted to the faculty of
The University of Utah
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

Department of Chemical Engineering

The University of Utah

August 2015

The University of Utah Graduate School

STATEMENT OF DISSERTATION APPROVAL

The dissertation of **Benjamin Brian Schroeder**

has been approved by the following supervisory committee members:

Sean T. Smith	, Chair	**06/10/15** Date Approved
Philip J. Smith	, Member	**06/10/15** Date Approved
David O. Lignell	, Member	**06/10/15** Date Approved
Jeremy Thornock	, Member	**06/10/15** Date Approved
Terry Ring	, Member	**06/10/15** Date Approved

and by **Milind Deo** , Chair/Dean of

the Department/College/School of **Chemical Engineering**

and by David B. Kieda, Dean of The Graduate School.

ABSTRACT

Scale-bridging models are created to capture desired characteristics of high-fidelity models within low-fidelity model-forms for the purpose of allowing models to function at required spacial and/or temporal scales. The development, analysis, and application of scale-bridging models will be the focus of this dissertation. The applications dictating scales herein are large-scale computational fluid dynamics codes. Three unique scale-bridging models will be presented. First, the development and validation of a multiple-polymorph, particle precipitation modeling framework for highly supersaturated $CaCO_3$ systems will be presented. This precipitation framework is validated against literature data, as well as explored for additional avenues of validation and potential future applications. Following this will be an introduction to the concepts of validation and uncertainty quantification and an approach for credible simulation development based upon those concepts. The credible simulation development approach is demonstrated through a spring-mass-damper pedagogical example. Bayesian statistical methods are commonly applied to validation and uncertainty quantification issues and the well-known Kennedy O'Hagan approach towards model-form uncertainty will be explored thoroughly using a chemical kinetics pedagogical example. Additional issues and ideas surrounding model-form uncertainty such as the identification problem will also be considered. Bayesian methods will then be applied towards the creation of a scale-bridging model for coal particle heat capacity and enthalpy modeling. Lastly, an alternative validation and uncertainty quantification technique, known as consistency testing, will be utilized to create a scale-bridging model for coal particle devolatilization. The credibility of the devolatilization scale-bridging model due to the model development process is assessed and found to have benefited from the use of validation and uncertainty quantification practices.

To my family for their unwavering support.

CONTENTS

LIST OF TABLES

ACKNOWLEDGMENTS

I would like to acknowledge the many people who made this dissertation possible. First, my advisor, Sean Smith, with whom I had the valuable opportunity to collaborate with and learn from throughout my graduate research. Phil Smith and Jeremy Thornock with whom I gained insight into application-based research and gained exposure to fascinating real world problems. Brigham Young University professors David Lignell and Thomas Fletcher and University of California - Berkeley professors Michael Frenklach and Andrew Packard with whom I had great collaborative experiences. Richard Hills, Kenneth Hu, and John Red-Horse who provided me helpful feedback and input during my summer at Sandia National Laboratories, as well as Walt Witkowski who was a great manager to work with and overall reference for the entire field of verification, validation, and uncertainty quantification. The staff at the Institute for Clean and Secure Energy who help me through the administrative process of being a graduate student. And last but not least Alex Abboud, Derrick Harris, Benjamin Isaac, and John Camilo Parra Alvarez with whom I had a great time conducting research with and going through graduate school.

CHAPTER 1

INTRODUCTION

1.1 Overview

Model building is a significant enterprise within simulation science. Models can be built to describe most multiphysics phenomena, albeit creation of a model does not guarantee satisfactory performance. A quote from George E. P. Box speaks towards this theme of model development: 'remember that all models are wrong; the practical question is how wrong do they have to be to not be useful' [16]. This dissertation will focus upon processes of building scale-bridging models. Such models are not meant provide new descriptions of physically phenomena, but capture and transfer desired portions of available descriptions, so that larger-scale models can couple a wide range of physical phenomena together in a feasible manner. The larger scale models for which models are developed herein are large-scale computational fluid dynamics codes. Scale-bridging models of $CaCO_3$ mineral, multiple polymorph, particle precipitation in highly supersaturated systems for carbon sequestration applications and coal particle heat capacity, enthalpy and devolatilization for oxy-fire coal boiler applications will be utilized to both demonstrate model building techniques as well as hold scientific value in the form to which they are developed. Validation and uncertainty quantification is a field of study dedicated towards increasing the value of simulation science. Methodologies from the validation and uncertainty quantification community will be applied towards testing and increasing credibility of scale-bridging models.

1.2 Scale-Bridging

Large-scale, multiphysics simulations utilize submodels to describe physical phenomena, but can not always afford high-fidelity submodels. Additionally, high-fidelity models may be inefficient or a waste of resources for the large-scale simulation because high-fidelity models will likely produce information that the simulation does not require. When developing physics models for implementation into simulations, the model must be designed to function at appropriate temporal and spacial scales. Scale-bridging models are created to translate

a model from higher-fidelity to lower-fidelity, so that the model functions at the proper scale for the application. Of course not all of the characteristics of the high-fidelity version of the model can be translated with the scale-bridging. The lower-fidelity version can be build based upon the ability to recreate the characteristics important to the application. Scale-bridging models differ from surrogate models in that they are physics based, albeit contain many approximations. An example of a simple scale-bridging model is a chemistry lookup-table tabulated across reaction extent, where the system contains many competing reactions.

Scale-bridging models present a unique opportunity for validation and uncertainty quantification methodologies. Regions of model validity for a scale-bridging model can be characterized based upon its high-fidelity equivalent. The uncertainty in the high-fidelity model can also be transferred to the scale-bridging model through this comparison, thus still allowing the uncertainty to be propagated forward through the model to the application predictions. Credibility of submodels is a concern for large-scale simulations, but if a scale-bridging model is developed through a validation and uncertainty quantification based methodology, the credibility of the high-fidelity model may become ingrained in the submodel. This should reduce uncertainty in application predictions.

1.3 Research Objectives

Objectives of this dissertation come at face value for the models developed, as well as in the processes utilized in their creation and improvement. Models created were developed with the intention of being utilized within computational fluid dynamics codes by other researchers postdevelopment. Given that a model created functioned for the desired application, understanding the model's validity, uncertainty, and credibility can be undertaken.

The first objective of this research was to develop a model capable of capturing the precipitation and mineral evolution processes for highly supersaturated $CaCO_3$ systems. Validation of this model was a requirement identified early in the model development process and thus the model development was largely led by the objective of fulfilling the necessary validation. The validation metric determined for the precipitation model was matching a set of experimental data extracted from the literature including multiple data types.

The next major objective was to develop a foundational basis for the utility of applying validation and uncertainty quantification towards the development of scale-bridging models. Validation and uncertainty methodologies were studied and a process for credible simulation

development is proposed. Model-form uncertainty is a significant issue within the development process for scale-bridging models and thus methods dealing with this uncertainty were required. Bayesian and engineering based methods of approaching such problems were explored.

The final objective was the creation of scale-bridging models where the process development was governed by validation and uncertainty quantification techniques. The application that drove the development of such scale-bridging models was oxy-fired coal boiler simulations. Heat capacity and enthalpy scale-bridging models for coal particles were created using Bayesian calibration, allowing greater transparency in the model's calibration and thus providing understanding of the model's validity and uncertainty. Coal particle devolatilization was another piece of physics that required scale-bridging. Devolatilization was known to contain significant amounts of uncertainty and thus an extensive effort towards the creation of a credible model was required. This objective included the creation of the scale-bridging model, as well as a demonstration of the improved product credibility produced utilizing validation and uncertainty quantification methods.

1.4 Organization of Dissertation

This dissertation is organized roughly in the order in which the research was undertaken. The exception to this is the philosophy within the process for developing credible simulations, which evolved throughout the time-frame of the entire research process. The two chapters focusing on particle precipitation are refined and expanded from a paper published on that work [127]. The chapter on coal devolatilization was also formulated with future publication planned. The credible simulation development discussion and application-problem walkthrough were created with the idea of being utilized in an upcoming course on verification, validation and uncertainty quantification being taught at the University of Utah, University of California - Berkley, and Brigham Young University during Fall 2015. Organization of this dissertation reflects its general theme: demonstrate scale-bridging model development and application, then introduce and explore concepts of validation and uncertainty quantification, and finally show that the integration of validation and uncertainty quantification methods into the process of creating scale-bridging models produces a superior product. Scale-bridging models can vary in range of physical complexity from single phenomena to many and strongly correlated physical phenomena. The first example provided is an instance of a scale-bridging model on the more complicated end of the spectrum.

A framework for modeling multiple-polymorph, particle precipitation processes within highly supersaturated $CaCO_3$ systems is proposed within Chapter 2. This framework description includes the approach taken towards the aqueous phase, utilizing population balance equations to describe the solid-phase evolution, physics submodels implemented within the population balances in order to capture the particle population's evolution, and a mixing model used to investigate mixing effects. A significant effort was taken towards modeling interfacial tension within the framework and this work is described last. All physical phenomena within the precipitation framework are formulated in manners allowing the physics to be coupled, which then permits comparison with experimental systems of interest.

An application study of the precipitation framework is presented within Chapter 3, where literature data are compared against framework predictions. A description of corrections made to the literature data due to discovered inconsistencies is first outlined. Next, a few of the potential configurations tested throughout the model development process are discussed. Following the configuration analysis, the final configuration selected for comparison with experimental data is presented and compared with the literature dataset. Additional outputs from the modeling framework are then demonstrated to display the potential utility of this modeling framework for similar applications. A brief illustration of the sensitivity to uncertainty in parameter values is shown in the last section of this chapter. This acknowledgment of uncertainty effects on the framework's predictions foreshadowed and motivated the application of means for increasing scale-bridging model's credibility.

Within Chapter 4 the ideas of validation and uncertainty are outlined and two methodologies for approaching model validation and uncertainty quantification are introduced and briefly described. The two methodologies considered are probabilistic Bayesian and consistency constraints. A process and philosophy for credible simulation design, containing applications for the validation and uncertainty quantification tools outlined, is presented. This design process is then demonstrated through a mass-spring-damper application where both Bayesian and consistency constraint approaches are used.

Continued exploration of Bayesian methodologies through considering a significant validation and uncertainty quantification issue, model-form uncertainty, is contained within Chapter 5. Overviews of Bayesian theory, Gaussian processes, and Markov Chain Monte Carlo methods are first introduced. Next, the particular Bayesian methodology known as the Kennedy O'Hagan approach is introduced and the algorithm utilized for implementing this approach within this chapter is outlined. Following this background information and

introductions is a pedagogical reaction-kinetics example further demonstrating the Kennedy O'Hagan methodology. A basic application to this pedagogical example is walked-through. Following the basic analysis is a description of a significant issue within model-form uncertainty problems known as the identification problem. Additional methods explored within the confines of the pedagogical example are then sketched out: Bayesian model comparison, multi-input Gaussian processes, and constrained Gaussian processes. Finally, forward propagation of the calibrated parameter distributions is shown.

Some of the same Bayesian techniques introduced and implemented within Chapter 5 are then applied towards parameter calibration of a scale-bridging model for coal particle heat capacity and enthalpy within Chapter 6. This problem acts as a simple demonstration of validation and uncertainty quantification based scale-bridging model development. First the application's motivation for this scale-bridging model is provided. Following this is the derivation of the model-form implemented and description of the creation of probabilistic characterizations used within the Bayesian calibration. Concluding this chapter are comparisons of a few model-form choices and demonstration of model performance. Although this example utilized Bayesian methods, consistency constraints could have been implemented in a similar manner, as is then shown through the next scale-bridging example.

Chapter 7 describes the development process utilized to create a scale-bridging model for coal particle devolatilization. First, the validation and uncertainty quantification based process that guided the model development is outlined. Next, the application space necessitating the scale-bridging model is presented. Following this is a brief description of the high-fidelity model utilized as the basis for the scale-bridging, as well as investigations into uncertainty generating components of the high-fidelity model. The model-form used as the reduced model for scale-bridging is then reported. Continuing along the model development's process-flow, the theory of consistency testing is presented and then the application of this methodology on the devolatilization scale-bridging model is analyzed. The last sections of this chapter discuss the credibility of the model developed and demonstrate how the validation and uncertainty quantification based model development process utilized can lead to continued model refinement.

CHAPTER 2

MULTIPLE-POLYMORPH PARTICLE PRECIPITATION

2.1 Introduction

Particle-precipitation modeling literature encompasses a wide range of physical phenomena including nucleation theory, growth mechanisms, dissolution mechanisms, aggregation kinetics, and Ostwald ripening or coarsening effects. For precipitating systems with multiple polymorphic forms, there may also be transitions from less thermodynamically stable polymorphic forms to more stable, or Ostwald's step rule, for the given system conditions [108]. Such transitions between polymorphic forms within the confines of experimental observations requires large shifts in the system's aqueous composition, or polymorphs' relative supersaturations. In order to capture the full evolution of precipitating systems, ionic aqueous-phase chemistry and solid-phase physical phenomena, whose effects often differ depending upon the portion of the particle population under consideration, must be coupled. Descriptions of the aqueous phase chemistry, solid-phase phenomena, and coupling of those phases is described in the following chapter. These physical descriptions are developed for the purpose of testing the coupled physics in a relatively simple reactor environment prior to use within large-scale computational fluid dynamics (CFD) simulations. Thus, the models must be developed at scales appropriate for the final application or a scale-bridging formulation. While many of the physical phenomena occurring with such precipitation systems could be described with computational chemistry, or other fine scale computations, here models must be developed at spacial scales suitable for the CFD discretization, while also being able to evolve at the temporal scale at which CFD operates.

Adapted with permission from B. Schroeder, D. Harris, S. Smith, and D. Lignell, *Theoretical framework for multiple-polymorph particle precipitation in highly supersaturated systems*, Cryst. Growth Des., 14 (2014), pp. 1756-1770.

2.2 Ionic Aqueous-Phase Chemistry

The ionic aqueous-phase chemistry of systems considered for calibrating and validating the precipitation framework will be considered to be at equilibrium, or assumed to be on far shorter timescales than all of the solid-phase physics. To model this equilibrium the open-source software toolkit Cantera [50] will be used. Cantera uses temperature dependent chemical kinetics and thermodynamics, as well as the Pitzer relations to calculate chemical activities of species in solution once the aqueous composition is known. Temperature-dependent solubility-product correlations for polymorphic forms of $CaCO_3$ can be found in Plummer [109] and Brecevic [19]. Using the ratio of the polymorph-specific, indexed j, equilibrium solubility-products $[CaCO_3]_{eq,j}$ and the activity of the aqueous-phase calcium carbonate $[CaCO_3]$, the respective supersaturation ratios, S_j, can be calculated. After comparing the experimental ion-activity product (IAP) results reported by Ogino et al. [106] with the equilibrium chemistry model of Plummer and Busenberg [109], which the aforementioned results were based upon, it was hypothesized that an error existed within the original code used to create those results. This hypothesis and the method of deriving the means of correcting for this error were presented within Schroeder et al. [127] and are described in Section 3.2.

2.3 Solid-Phase Particle-Phenomena Models

2.3.1 Population Balances

In order to track the evolution of the solid-phase materials, population balance equations (PBEs) will be utilized. Randolph and Larson [114] popularized the use of PBEs in particle science and Ramkrishna [113] is known for presenting greater mathematical characterization of the use of such methods. Population balances are commonly utilized in the particle science community and have been shown to accurately track crystal populations for systems similar to those of interest [133, 102, 103]. Population balance equations track the evolution of distributions describing population's physical characteristics such as size, shape, densities, etc. The following equation is the general form of the population balance equations that will be utilized to track the particle populations of each polymorph, where the particle characteristic or internal coordinate being transported is the particle radial size, r,

$$\frac{\partial \eta}{\partial t} + \frac{\partial}{\partial r}[G(r)\eta] = B_N(r) + A(r). \tag{2.1}$$

Here G is the growth rate, B is the birth rate, and A is the aggregation rate. No external coordinates will be tracked explicitly by the PBEs, effectively treating the system as a single well-mixed batch-reactor. A mixing model will be utilized to explore basic mixing

effects upon the systems of interest. Another simplifying assumption utilized, is to omit any cross-polymorph aggregation. This assumption is justified by aggregation being a second order effect for the systems of interest.

Instead of incurring the computational cost of transporting the entire radial distribution, the method of moments will be utilized to represent the distribution by a few low order moments. By defining the α^{th} radial moment of the distribution, η, as $m_\alpha \equiv \int_0^\infty r^\alpha \eta \, \mathrm{d}r$, Eq. (2.1) can be written as

$$\frac{\partial m_\alpha}{\partial t} - \alpha \int_{r_c}^{\infty} r^{\alpha-1} G(r) \eta \, \mathrm{d}r = \int_{r_c}^{\infty} r^\alpha B_{\mathrm{N}}(r) \, \mathrm{d}r + \int_{r_c}^{\infty} r^\alpha A(r) \, \mathrm{d}r. \tag{2.2}$$

A closure technique is needed for Eq. (2.2) due to the integration of the growth and aggregation terms that depend upon the unknown distribution. Quadrature based methods of solving PBEs, such as the quadrature method of moments (QMOM), were introduced by McGraw [94] as a means of providing closure to the aforementioned integrals, and large amounts of progress in the development of such methods has been published by Dr. Rodney Fox's group at Iowa State University [39, 92, 93, 91]. QMOM approximates integrals using Gaussian-quadrature with its associated weights, w_k, and abscissae, R_k, for each quadrature node $k = 1, 2, ..., 2N-1$, or defined such that lower-order moments are satisfied as

$$m_\alpha = \sum_{k=1}^{N} w_k R_k^\alpha, \quad \text{for } \alpha = 0, 1, \ldots, 2N-1. \tag{2.3}$$

The product-difference algorithm, as derived by Gordon [51], can be used to solve for the weights and abscissae given the moments. Thus the QMOM form of Eq. (2.1) is

$$\frac{\mathrm{d}m_\alpha}{\mathrm{d}t} - \alpha \sum_{k=1}^{N} w_k R_k^{\alpha-1} G(R_k) = B_{\mathrm{N},\alpha} + A_\alpha. \tag{2.4}$$

Alternatively, the weights and abscissae can be directly transported, which is known as the direct quadrature method of moments (DQMOM) [90]

$$\sum_{k=1}^{N} \left[R_k^\alpha \frac{\mathrm{d}w_k}{\mathrm{d}t} \right] - \alpha \sum_{k=1}^{N} w_k R_k^{\alpha-1} \left[G(R_k) + \frac{\mathrm{d}R_k}{\mathrm{d}t} \right] = B_{\mathrm{N},\alpha} + A_\alpha. \tag{2.5}$$

For the current framework implementation, the DQMOM form of the PBEs will be solved. The birth

$$B_{\mathrm{N},\alpha} \equiv \int_{r_c}^{\infty} r^\alpha B(r) \, \mathrm{d}r \tag{2.6}$$

and aggregation-rate integrals

$$A_\alpha \equiv \int_{r_c}^{\infty} r^\alpha A(r) \, \mathrm{d}r \tag{2.7}$$

both need means of solution and will be discussed within the the following sections along with growth mechanisms that are specific to both polymorph and system conditions.

2.3.2 Nucleation

The birth source term, B, for this application will be modeled using a classical homogeneous-nucleation mechanism [71]

$$J = zk_f C(1)C_e(i_c), \tag{2.8}$$

where z is the Zeldovich factor, k_f is the forward reaction-rate coefficient for molecular growth of a cluster, $C(1)$ is the number density of single precipitant molecules in solution, and $C_e(i_c)$ is the equilibrium-based number density of particles of the size currently nucleating. The particle or embryonic-cluster size number-density is calculated with the Boltzmann equation

$$C_{\mathrm{e}}(i) = C(1)\exp^{\frac{-\Delta G(i)}{k_{\mathrm{B}}T}}, \tag{2.9}$$

where $\Delta G(i)$ is the Gibbs free energy, k_B is the Boltzmann constant, and T is the temperature. Due to nucleation causing the system to deviate from equilibrium, the Zeldovich factor is used to correct for the equilibrium assumption upon which the Boltzmann distribution is based [149]

$$z \approx \left(\frac{(\frac{-d^2G}{di^2})_{i=i_C}}{2\pi k_{\mathrm{B}}T}\right)^{1/2} \approx \left(\frac{\Delta G_C}{3\pi k_{\mathrm{B}}Ti_C^2}\right)^{1/2}. \tag{2.10}$$

A variety of forward reaction-rate coefficients exist, but the two forms explored throughout this research and commonly found in the literature are interface-transfer-limited ($k_{f,i}$) and diffusion limited ($k_{f,d}$) varieties [71]

$$k_{f,i} = D(6\pi^2\nu)^{1/3}i^{2/3} \tag{2.11}$$

$$k_{f,d} = D(48\pi^2\nu i)^{1/3}. \tag{2.12}$$

Here ν is the molecular volume, D is the diffusion coefficient, and i is the number of molecules composing the cluster. Any of these equations can be easily converted into terms of cluster/particle radius by assuming spherical shape ($i_c = \frac{4\pi}{3\nu}r_c^3$).

The birth-rate due to nucleation can be mathematically described by a delta function at the desired size, r_c, multiplied by the nucleation rate or $B_N = \delta(r - r_c)J$. Thus the birth-rate integral for Eq. (2.4) will be

$$B_{N,\alpha} \equiv \int_{r_c}^{\infty} r^{\alpha}B_N(r)\,\mathrm{d}r = r_c^{\alpha}J. \tag{2.13}$$

This equation is shown in terms of radial-size because the PBEs in which it is utilized are based upon transporting particle radial-size characteristics. The size at which particles are

inserted into the PSD will be the critical radii or the size at which it becomes energetically favorable for particle embryos to nucleate as is defined by the Gibbs free energy equation for molecular clusters

$$\Delta G(i) = -k_B T \ln(S)(i-1) + (36\pi\nu^2)^{1/3}\sigma(i^{2/3}-1). \tag{2.14}$$

Here σ is the interfacial tension. Eq. (2.14) has been written relative to a single molecule to maintain self consistency [70, 47]. Solving the derivative of Eq. (2.14) with respect to radius or number of molecules and then setting equal to zero corresponds to the point at which the energetics of additional particle volume overtake the surface energy's resistance, or the critical size

$$i_C = \frac{32\pi\nu^2\sigma^3}{3(k_B T \ln S)^3} \qquad \text{or} \qquad r_C = \frac{2\sigma\nu}{k_B T \ln S}. \tag{2.15}$$

As the system's composition changes, so too will the rate of nucleation and the size at which particles become energetically favorable or enter the PBE. This is reflected within the integral limits of Eq. (2.13) that can now be understood to change over time as the system conditions vary.

While the above equations are generally based upon static interfacial tension values, the use of size dependent interfacial tension, as is discussed further in Section 2.3.7, provides further complications. The previously derived critical Gibbs free energy was based upon solving $\frac{d\Delta G(i)}{di} = 0$, which can be rederived once the size dependent form of interfacial tension has been chosen

$$0 = \frac{k_B T \ln S}{\nu} r_C^2 - 2\sigma_\infty r_C + 2\sigma_\infty \delta_T. \tag{2.16}$$

When solving it was found that taking the larger root kept the system physical and the actual critical size utilized was 1% larger in order to numerically separate where particle nucleated and growth began.

The nucleation rate, J, has thus far been treated as a quasisteady state model, where it has been assumed that the prenucleation embryonic-cluster distribution is fully established. For systems infinitely fast mixing, this assumption would prove false, but for the current system with finite mixing rate its validity is uncertain. As the Ca_2^+ and CO_3^{2-} solutions are mixed, $CaCO_3$ molecules will form and a distribution of embryonic clusters of the $CaCO_3$ molecules will evolve towards the quasisteady state Boltzmann distribution. During this period of transition, the embryo distribution will have greater weight on smaller sized clusters. During the transition towards a quasisteady embryo distribution, nucleation can be

described by transient nucleation models. Such models utilize the transitional time frame τ, known as the time lag or induction time, as the timescale upon which transitioning to steady-state nucleation is scaled. The most well known model describing transient nucleation is that by Kashchiev [69]

$$J(t) = J\left[1 + 2\sum_{m=1}^{\infty}(-1)^m \exp(-m^2\frac{t}{\tau})\right], \tag{2.17}$$

where $\tau = 4/z^2\pi^3 k_f N(1)$ [73].

Alternative nucleation mechanisms exist and could potentially be present in the conditions of the systems of interest. A secondary heterogeneous-nucleation mechanism [71] was implemented during framework development. Secondary heterogeneous nucleation allows for nucleation on particles already present in the system, as opposed to primary heterogeneous nucleation that would occur on system boundaries and impurities. Using this heterogeneous nucleation mechanism introduces additional unknown model parameters such as: the contact angle at which one polymorph nucleates on particles of each of the other polymorphic forms, and the density of nucleation sites for each polymorphic form. These additional unknown parameters had to be fitted. Within recent literature it has been suggested that ACC might follow a nonclassical nucleation pathway [118, 42]. Although qualitatively ACC's nonclassical nucleation has been described, no mathematical equivalent has been proposed thus far. The use of classical nucleation mechanisms to describe ACC's appearance were investigated by Harris (2013) [60] and found to perform in a satisfactory manner.

2.3.3 Growth

Growth mechanisms can differ for each polymorph due to differing crystal structures and even then the growth mechanism will change as the system's supersaturation values relative to each polymorph also evolve. The following growth mechanisms were selected based upon the system conditions of the validation system, but it is recognized that different mechanisms would be appropriate if the system's initial conditions were altered and that the chosen mechanisms are rough approximations. Because the validation systems utilized have high initial supersaturation values, diffusion-limited growth [100] (Eq. (2.18)) will be the controlling mechanism for most of the crystalline polymorphs until their respective supersaturations have been greatly diminished [27]. The diffusion-limited growth mechanism, assuming equilibrium boundary conditions, can be expressed as

$$G(r) = \frac{D}{\rho r}[\mathrm{CaCO_3}]_{eq}(S - \bar{S}), \tag{2.18}$$

where D is the diffusion coefficient, ρ is the polymorph's molar density, and $\bar{S}$ is a ratio of the activities accounting for the difference between a finite-radius particle and an infinite flat surface [100]. This activity ratio is often set to unity, likely due to describing geological precipitation instead of particles, but in order to capture coarsening effects it can be derived from the Kelvin equation to be $\bar{S} \equiv e^{2\sigma/\rho RTr}$ [27, 102] or $\bar{S} \equiv e^{2\sigma_\infty(1-\delta/r)/\rho RTr}$ once size dependent interfacial tension is included, Eq. (2.57). Eq. (2.18) assumes spherical particles. While this assumption should be valid for ACC, vaterite, and calcite, aragonite is known to take geometries better described as cylindrical. A diffusion-limited growth model for cylindrical particles was then derived as

$$G(r) = \frac{7}{6\ln(2)}\frac{D}{\rho r}[\mathrm{CaCO_3}]_{\mathrm{eq}}(S - \bar{S}), \tag{2.19}$$

where it was assumed that the diameter-to-height aspect ratio was 1:6 and the concentration boundary layer of the particles was within an order of the particle size.

Once respective supersaturation values for each polymorph become close to $\bar{S}$, the growth rates are known to be determined by surface-controlled mechanisms, such as this singular-sourced screw-dislocation mechanism

$$G(r) = K_r(S - \bar{S})^2, \tag{2.20}$$

where K_r is an empirical reaction-rate constant [76, 75, 3, 13]. The recently published surface-reaction limited mechanism for calcite from [145] that accounts for pH and ion ratio effects will be utilized. Although ACC has supersaturation values relatively lower than the other polymorphic forms, diffusion-limited growth will still be utilized due to ACC's minimal crystal structure.

2.3.4 Dissolution and Death

Mathematically, dissolution appears as a negative growth term and complete particle dissolution is implemented as a death term or negative birth event. As the system's supersaturation relative to each polymorph reaches an undersaturated state, the system displays Ostwald's step rule [108], or the gradual descent through increasingly more stable mineral forms until reaching the thermodynamically most stable polymorph. Dissolution, as it is defined within this framework, begins once a portion of a polymorph's PSD becomes undersaturated or $S < \bar{S}$. It is through this definition of dissolution that coarsening or Ostwald ripening [81] within a polymorph's PBE can be captured. Without the inclusion of $\bar{S}$, all sizes of a mineral form would begin dissolution simultaneously.

There is currently ambiguity in ACC's dissolution mechanism [118, 15], so surface- and diffusion-limited mechanisms, Eq. (2.20) and (2.18), with fitted rate constants were investigated. Due to the dissolution mechanisms of vaterite and aragonite having not been extensively studied across the composition and temperature ranges of interest, the diffusion-limited mechanism (Eq. (2.18)) was implemented as a general estimate. Growth rates of more stable polymorphs have been found to be a limiting factor within $CaCO_3$ systems [118, 77, 75], so the use of diffusion-limited dissolution for these two metastable polymorphs should allow the growth mechanisms of more thermodynamically favorable mineral forms to dictate the system dynamics. Although calcite's dissolution has a minor appearance/effect upon the validation systems, calcite's dissolution mechanisms are well characterized across a wide range of compositions and temperatures [22, 58, 24]. The surface-limited reaction mechanism of Plummer et al. (1978) [110] will be used to model calcite's dissolution due to the mechanisms functionality across composition and temperature.

In order to remove particles from the PSD, a left boundary must be selected, r_{cutoff}, and a death term D must be specified. A minimal distribution weight at which a distribution environment will no longer have death enacted upon it will also be set for numerical reasons. The chosen left boundary must have a negligible effect upon the equilibrium-chemistry, but is still necessary in order to stop the PSD from encountering numerical issues. Because this death model serves as numerical convenience, but does not represent a physical process, an empirically defined equation is utilized

$$D_i = \frac{-30kw_i}{\Delta t}, \tag{2.21}$$

where

$$k = 0.25\left(1 - \text{erf}\left(8\ln(\tfrac{R_i}{r_{\text{cutoff}}})\right)\right)\left(1 - \text{erf}\left(5\log_{10}(\tfrac{10^{-2.5}}{w_i})\right)\right). \tag{2.22}$$

A death term such as this was mentioned within Yuan (2012) [148], but the exact form was developed for this specific application. The criteria of selecting when death occurs in the system, or left boundary, were set at a particle radial size of $6 \times 10^{-2}\mu m$ with a respective distribution weight of $10^{-3}\#/m^3$.

2.3.5 Aggregation

Although aggregation is not believed to be an effect of primary importance within the systems of interest and is not explicitly measured through the experimental data available, it is included for completeness and future flexibility in the theoretical description provided by the framework. Aggregation is the processes by which two individual particles collide

and form a singular particle. Within the context of PBEs, a particle of the combined size is introduced into the population while two particles of the original sizes are lost. Mathematically this requires two source terms for the PBE, a birth and a death term. A PBE that only considers these two aggregation source terms is known as the general aggregation equation

$$\frac{\partial n'(v;t)}{\partial t} = \frac{1}{2}\int_0^v \alpha'(v-v',v')n'(v-v';t)n'(v';t)\,\mathrm{d}v' - n'(v;t)\int_0^\infty \alpha'(v,v')n'(v';t)\,\mathrm{d}v', \tag{2.23}$$

where n' is the number density of particles on a volume basis, α' is the aggregation kernel on a volume basis, and v and v' are the volumes of the two particles forming the new aggregate. The first term of the right hand side of Eq. (2.23) represents the creation of a new particle of size v, while the death of particles due to the aggregation is captured by the second term. Following the derivation by Marchisio et al. (2003) [92], this aggregation based population balance can be converted into its equivalent form in terms of weights and abscissae that can be utilized in the current framework

$$A_\alpha \approx \frac{1}{2}\sum_{i=1}^{N} w_i \sum_{j=1}^{N} \beta(R_i,R_j)w_j(R_i^3+R_j^3)^{\frac{\alpha}{3}} - \sum_{i=1}^{N} R_i^\alpha w_i \sum_{j=1}^{N}\beta(R_i,R_j)w_j. \tag{2.24}$$

The radial-based aggregation kernel β is the product of the radial-based collision-frequency kernel β^* and the collision efficiency Ψ [80],

$$\beta(r_i,r_j) = \beta^*(r_i,r_j)\Psi. \tag{2.25}$$

While there exist many collision-frequency kernels, a Brownian motion frequency kernel [30]

$$\beta_{\text{Brownian}} = \frac{2k_BT}{3\mu}\frac{(r+r')^2}{rr'}, \tag{2.26}$$

will be implemented due to the large quantities of small particles within the systems. The collision efficiency ranges from zero to one reflecting the probability of an aggregation event given that a collision has occurred. To approximate this probability a semi-empirical model that balances the bonding forces/growth between two particles against tensile and shear stresses of the fluid dynamics will be implemented [64, 80, 4].

Collision efficiency must be considered due to the fact that not all collisions result in aggregation. Fundamentally the collision efficiency can be thought of as the probability that forces bonding the particles together overcome forces pulling the particles apart. At low ionic strengths the balance between attractive and repulsive forces influences the efficiency,

but at high ionic strengths the repulsion greatly diminishes and the efficiency goes towards one [92]. While modeling this factor is an area of current research, one semi-empirical method introduced by Liew et al. (2003) [80] has been explored in recent literature [26, 65] and will be used for this application. This model creates a dimensionless strength variable (M) out of the factors effecting the likelihood of the formation of a new particle when two particles collide. Dimensionless strength is derived by comparing the strength of the bond formed between the two particles (σ) with the force from the surrounding fluid upon that bond (F/A),

$$M = \frac{\sigma}{F/A}. \tag{2.27}$$

The force placed on the bond is derived from hydrodynamic theory to be $F \propto \mu\dot{\gamma}d^2$, where μ is the fluid viscosity, $\dot{\gamma}$ is the shear rate, and d is the particle diameter. Previous versions of this model assumed that collisions occurred at single points, but this version allows for collisions to occur along a line. This seems reasonable due to crystal structures generally being composed of faces and edges rather than points. The area is set to $A = L(\frac{G}{\dot{\gamma}\sin(\Omega)})$, where G is the growth rate and Ω is a shape factor. The growth rate is included because it describes that rate of enlargement of the bond. Using the defined terms in Eq. (2.27) along with assuming $\mu\dot{\gamma}^2 \propto \rho\bar{\epsilon}$, dimensionless strength can be expressed as

$$M = \frac{L\sigma^* G}{\rho\bar{d}^2\bar{\epsilon}}, \tag{2.28}$$

where ρ is the fluid density, $\bar{\epsilon}$ is the mean turbulent dissipation energy, $\sigma^* = \sigma/\sin(\Omega)$, and $\bar{d}$ is the average particle diameter. For the average particle diameter a geometric average of the two particles will be used. If the particle growth rate being used is size dependent, it seems reasonable to base the growth rate on the larger of the two particles colliding because this value will be closer to the final particle size. Also, if a moment method is being used, the growth rate has likely already been solved for the particle sizes colliding, but not the new particle size. Because it was known that $M \propto \Psi$ at low efficiency and goes to one at high inefficiency, Ψ takes the following functional form

$$\Psi = \frac{M/M_{50}}{1 + M/M_{50}}, \tag{2.29}$$

where M_{50} is the M value for fifty percent efficiency. Values for the variables M_{50}, $\sigma*$, and L can be found in the literature for minerals such as calcite and vaterite, as can be seen in Table 2.1. Due to not being found in the literature, values of 0.1 and 0.5 $N\ m^{-1}$ were approximated for ACC and aragonite based upon the values for the other polymorphs

Table 2.1: Example of values for Eq. (2.28) from the literature.

—	Calcite [80]	Vaterite [4]
$L\sigma^*/M_{50}$ $[N\,m^{-1}]$	0.72 ± 0.03	0.18 ± 0.02

shown in Table 2.1. For this model input values used included a fluid density of 998.2 $kg\ m^{-3}$, a fluid dynamic viscosity of $1.002 \times 10^{-3}\ kg\ s^{-1}m^{-1}$, and a turbulent dissipation energy of 5000 $J\ kg$.

2.3.6 Mixing

In order to perform a preliminary evaluation of mixing effects upon the systems evolved within the precipitation framework, a simplistic mixing model will be included. The mixing model selected for this task is a three-environment, multienvironment micromixing model, where initially two environments represent the two feed streams into the precipitation system [39, 143]. As the two feed streams begin to mix together a third environment emerges into which the other two environments are redistributed over time. It is within this middle or third environment that all of the particle physics will then occur. This mixing model effectively scales the PBEs by the weight of the third mixing environment. The two initial environments shrink in size, but maintain constant compositions. Within [143] it was shown that this type of three-environment, multienvironment mixing model performed well in capturing low order moment characteristics of precipitation systems such as mean particle size and mean particle number density. The amount of mixing was calculated as

$$\text{Mixing} = -\frac{1}{t_{mix}} \cdot \frac{Z'^2}{Z(1-Z) - Z'^2}, \tag{2.30}$$

where t_{mix} is the mixing time, Z is the mean mixture fraction, and Z'^2 is the mixture fraction variance. Mixing is then incorporated into the PBE as

$$\sum_{k=1}^{N} \left[R_k^{\alpha} \left(\frac{\mathrm{d}w_k}{\mathrm{d}t} + w_k(\text{Mixing}) \right) \right] - \alpha \sum_{k=1}^{N} w_k R_k^{\alpha-1} \left[G(R_k) + \frac{\mathrm{d}R_k}{\mathrm{d}t} \right] = B_{N,\alpha} + A_{\alpha}. \tag{2.31}$$

2.3.7 Interfacial Tension

Interfacial tension will be a key parameter in the nucleation calculations and major contributor to the uncertainty within the framework. Common methods of determining interfacial tension from experimental data include utilizing the Ostwald-Freundlich relation with measured solubility and nucleation rate/induction-period data for correlations [146]. Although interfacial tension values for many of the polymorphs of interest are extensively

reported throughout the literature, as can be seen in Table 2.2, there is a wide variance in the values reported likely due to differing means of experimental determination. While interfacial tensions are dependent upon crystal structure, changing between polymorphic forms, they are also known to be dependent upon the system temperature and composition. Such temperature and composition functionality must be included within the precipitation framework in order to best capture the experimental data. Thermodynamically derived equations presented by van Oss [142] and Mersmann [96] allow for temperature and system composition dependence, but also involve more parameters that must either be experimentally determined or utilized as uncertain parameters for calibration.

2.3.7.1 Thermodynamic Relation Derivation

The following derivation of thermodynamic relations describing interfacial tension will stem off of the work of Mersmann (1990) [96]. The equations proposed by Mersmann were meant to be simplistic and dependent upon composition but not temperature. In order to create an interfacial tension model that included both dependencies, the Mersmann derivation can be altered. Within his derivation, Mersmann took the Guggenheim philosophical approach towards interfacial tension [54]. Where Gibbs (1928) [45] assigned interfacial tension to a purely mathematical plane between the two bulk regions, Guggenheim's theory allows for an interfacial region. Within Guggenheim's interfacial region lies all of the inhomogeneities associated with changing from one homogeneous phase to another. Having this transitional region exist allows for physical properties to be assigned to it as though it were another bulk phase. With Gibbs' approach, the placement of the mathematical plane between the the two homogeneous bulk phases effects the values of the properties assigned to the plane.

Mersmann started his derivation with the Gibbs isotherm equation

$$\sum n_i^p \mathrm{d}\mu_i + A\mathrm{d}\sigma = 0, \tag{2.32}$$

which can be seen to be an isothermal and isobaric version of the Gibbs adsorption equation

Table 2.2: A review of values for interfacial tension adapted from Donnet et al. (2009) [28].

Polymorph	σ[mJ m^{-2}]	Source
ACC	–	Not well established
Vaterite	$6.8 - 133$	[130, 89, 76, 48, 28]
Aragonite	150	[130]
Calcite	$19.5 - 280$	[23, 130, 52, 63, 48]

$$\mathrm{d}\sigma = -S_a \mathrm{d}T - \sum \Gamma_i \mathrm{d}\mu_i + V_a \mathrm{d}P. \tag{2.33}$$

Here n_i^p is the number of moles of component i in the interfacial phase p, μ is the chemical potential, A is the interfacial area, $\Gamma_i = n_i^p/A$ is the surface concentration, $S_a = S^p/A$ is the surface entropy, P is the pressure, and V_a is the measure of the interface's width. The Gibbsian framework utilizes excess properties such as the surface excess entropy, S_a^E, instead of their counterpart properties described in Eq. (2.33) due to the lack of an interfacial phase. To derive the equivalent of the Gibbs adsorption equation within a Guggenheim framework, the definition of the differential of internal energy for the interfacial phase and the general definition of internal energy can be utilized.

$$\mathrm{d}U = T\mathrm{d}S - P\mathrm{d}V + \sigma \mathrm{d}A + \sum \mu_i \mathrm{d}n_i \tag{2.34}$$

$$U = TS - PV + \sigma A + \sum \mu_i n_i. \tag{2.35}$$

Thus in a variety of the Gibbs-Duhem equation

$$0 = S\mathrm{d}T - V\mathrm{d}P + A\mathrm{d}\sigma + \sum n_i \mathrm{d}\mu_i \tag{2.36}$$

$$\mathrm{d}\sigma = -S_a \mathrm{d}T - \sum \Gamma_i \mathrm{d}\mu_i + V_a \mathrm{d}P.$$

From Eq. (2.33) an interfacial tension model with linear temperature dependence can be explored by assuming that the effect of pressure and chemical potential changes are less significant,

$$\mathrm{d}\sigma = -S_a \mathrm{d}T \tag{2.37}$$

$$\int_{\sigma(T_o)}^{\sigma(T)} \mathrm{d}\sigma = \int_{T_o}^{T} -S_a \mathrm{d}T \tag{2.38}$$

$$\sigma(T) = \sigma(T_o) - S_a(T - T_o). \tag{2.39}$$

Although in the initial compositions the system's of interest will be known and fixed, those compositions will change over time as precipitation occurs unless additional material is added to the system. Thus if Eq. (2.39) were implemented within the framework, it would likely be representing a weighted average interfacial tension that gives greater weight to values during the nucleation processes if it is calibrated against experiments with temporally varying compositions. This weighing would be due to the physical effects of interfacial tension being most evident during nucleation. The temperature dependent interfacial tension model performed well for the fixed initial composition experiments of Ogino et al. [106] it was compared with, but in order to be able to capture composition varying experiments, we need to begin again with Eq. (2.33).

In order to maintain both temperature and composition dependence, the Gibbs adsorption equation will be carried through Mersmann's derivation instead of the Gibbs isotherm equation. It will still be assumed that the system is isobaric,

$$\mathrm{d}\sigma = -S_a \mathrm{d}T - \sum \Gamma_i \mathrm{d}\mu_i. \tag{2.40}$$

Now the surface concentration (Γ) must be described. Guggenheim described the surface concentration as 'the number of moles of the component in a unit area of the surface layer,' and Mersmann interpreted this to mean the concentration of the component in a monolayer at the interface of the interfacial region and the solid-phase. In the Gibbsian framework, this term describes excess concentration of the component within the mathematical plane (or $n_{\text{total}} - n_{\text{phase 1}} - n_{\text{phase 2}} = n_\sigma$). To quantitatively describe the surface concentration, Mersmann assumes the particles in the monolayer are spheres and relates their surface area to their volume,

$$O_{m,Total} = \pi d_m^2 \qquad V_m = \pi d_m^3/6 \tag{2.41}$$

$$\rightarrow \; O_{m,Total} = \pi^{1/3} 6^{2/3} V_m^{2/3}. \tag{2.42}$$

The surface concentration is then described as $\Gamma_i = 1/(N_A O_{m,i})$, where Mersmann assumes that the molecules in the monolayer are 50% exposed to the liquid phase and N_A is the Avogadro constant. The molar volume is described in terms of the component's density in the solid-phase (c_i^s) or $O_{m,i} = \frac{\pi^{1/3} 6^{2/3}}{2} (c_i^s N_A)^{-2/3}$. Substituting this back into Eq. (2.40) yields

$$\mathrm{d}\sigma = -S_a \mathrm{d}T - \frac{0.414}{N_A} \sum (c_i^s N_A)^{2/3} \mathrm{d}\mu_i. \tag{2.43}$$

Using other geometric and coverage assumptions leads to different values for the surface concentration term. If it is assumed that water does not exist in significant quantities within the crystal lattice and that the surface of tension is located so that $\Gamma_{\mathrm{H_2O}} = 0$, the summation can be dropped and only the terms involving $\mathrm{CaCO_3}$ remain,

$$\mathrm{d}\sigma = -S_a \mathrm{d}T - \frac{0.414}{N_A} (c^s N_A)^{2/3} \mathrm{d}\mu. \tag{2.44}$$

Now an integration over Eq. (2.44) from a reference interfacial tension value at a reference chemical potential and temperature yields

$$\int_{\sigma(T_\circ, \mu^\circ)}^{\sigma} \mathrm{d}\sigma = -S_a \int_{T_\circ}^{T} \mathrm{d}T - \frac{0.414}{N_A} (c^s N_A)^{2/3} \int_{\mu^\circ}^{\mu} \mathrm{d}\mu \tag{2.45}$$

$$\sigma = \sigma(T_\circ, \mu^\circ) - S_a (T - T_\circ) - \frac{0.414}{N_A} (c^s N_A)^{2/3} (\mu - \mu^\circ). \tag{2.46}$$

The chemical potential can then be expressed in terms of activity [134, 10],

$$\mu_i = \mu_i^\circ + RT \ln a_i \tag{2.47}$$

$$\mu_{\mathrm{CaCO_{3,aq}}} = (1)\mu_{\mathrm{Ca^{2+}}} + (1)\mu_{\mathrm{CO_3{}^{2-}}} \tag{2.48}$$

$$= \mu^\circ_{\mathrm{Ca^{2+}}} + RT \ln a_{\mathrm{Ca^{2+}}} + \mu^\circ_{\mathrm{CO_3{}^{2+}}} + RT \ln a_{\mathrm{CO_3{}^{2-}}} \tag{2.49}$$

$$\mu^\circ_{\mathrm{CaCO_3}} = (1)\mu^\circ_{\mathrm{Ca^{2+}}} + (1)\mu^\circ_{\mathrm{CO_3{}^{2-}}} \tag{2.50}$$

$$\mu_{\mathrm{CaCO_{3,aq}}} = \mu^\circ_{\mathrm{CaCO_3}} + RT \ln(a_{\mathrm{Ca^{2+}}} a_{\mathrm{CO_3{}^{2-}}}). \tag{2.51}$$

In the equations above (1) is displayed as a reminder of the chemical complex's stoichiometry that must be included. Although it is a common practice to use the activities of the ions in solution to describe the chemical potential of the liquid complex, in order to maintain consistency with the supersaturation definition being utilized in our mineralization framework, the complexed form ($\mathrm{CaCO_{3,aq}}$) will be utilized to describe the chemical potential of interest. The reference chemical potentials cancel out. The equilibrium between the chemical potential of the sum of the ions and the complexed form can be utilized to replace the ionic activities,

$$K = \frac{a_{\mathrm{CaCO_{3,aq}}}}{a_{\mathrm{Ca^{2+}}} a_{\mathrm{CO_3{}^{2-}}}}. \tag{2.52}$$

Then the supersaturation ($\mathtt{S}$) and the liquid-solid complex equilibrium constant ($[\mathrm{CaCO_3}]_{eq}$) can be used to describe the complexed form activity,

$$\begin{aligned}
\sigma = \sigma(T_\circ, \mu^\circ) &- S_a(T - T_\circ) \\
&- \tfrac{0.414}{N_A}(c^s N_A)^{2/3}\big(\mu^\circ_{\mathrm{CaCO_3}} + RT \ln(a_{\mathrm{Ca^{2+}}} a_{\mathrm{CO_3{}^{2-}}}) - \mu^\circ_{\mathrm{CaCO_3}}\big) \\
= &\quad - 0.414 k_B T (c^s N_A)^{2/3} \ln(a_{\mathrm{Ca^{2+}}} a_{\mathrm{CO_3{}^{2-}}}) \\
= &\quad - 0.414 k_B T (c^s N_A)^{2/3} \ln\big(\tfrac{a_{\mathrm{CaCO_{3,aq}}}}{K}\big) \\
= &\quad - 0.414 k_B T (c^s N_A)^{2/3} \ln\big(\tfrac{\mathtt{S}[\mathrm{CaCO_3}]_{eq}}{K}\big) \\
\sigma(T, \mu) = \sigma(T_\circ, \mu^\circ) &- S_a(T - T_\circ) - 0.414 k_B T (c^s N_A)^{2/3} \ln\big(\tfrac{\mathtt{S}[\mathrm{CaCO_3}]_{eq}}{K}\big).
\end{aligned} \tag{2.53}$$

Which for each polymorph then takes the form

$$\sigma_j(T, \mu) = \sigma_j(T_\circ, \mu^\circ) - S_{a,j}(T - T_\circ) - 0.414 k_B T (c^s_j N_A)^{2/3} \ln\big(\tfrac{\mathtt{S}_j[\mathrm{CaCO_3}]_{j,eq}}{K}\big). \tag{2.54}$$

The reference interfacial tension value in Eq. (2.54) does not have as physical of a meaning as the reference in Eq. (2.39). Where in Eq. (2.39) the reference value was the actual interfacial tension value utilized at the reference temperature, within Eq. (2.54) the composition dependent term is not negligible at the reference temperature unless the current

activity coincides with the solid-liquid equilibrium value's. The reference interfacial tension value can still be used as a calibration parameter along with the surface entropy, leaving two free parameters per polymorphic form.

2.3.7.2 Size Dependence

Like most physical material properties, the definition of interfacial tension and its quantitative value begin to lose validity when looking at small amounts of material. At high supersaturation values, particles can nucleate at sizes where continuum approximations are questionable. Although it seems appropriate to use the interfacial tension of an infinity sized interface for particles that are millimeter and centimeter sized (known as the capillary approximation), the very definition and applicability of interfacial tension becomes questionable for nanoscale particles. With the recent expansion of nanoscale science, there have been many suggestions on how to scale properties with particle size [56, 55, 57]. This same question, specifically for interfacial tension, was addressed long ago by Tolman (1948) [137].

Tolman's model for radially dependent interfacial tension,

$$\sigma(r) = \sigma_{\text{bulk}}(1 + 2\delta_T/r)^{-1}, \tag{2.55}$$

is well established in the literature, but introduces a Tolman length term δ_T that is not well documented for solid particulates. Although there has been a great breadth of research into the Tolman lengths for liquid droplets, there has yet to be a consensus on either the magnitude or sign [138, 141, 97]. Following Kalikmanov (1997) [68], the Tolman length will be estimated as 20% of the radial size of a single molecule.

Within the derivation of Tolman's equation, it was assumed that the Tolman length is not a function of the particle size. Thus the presented equation is a truncation due to the unknown functionality of the Tolman length with regards to the particle size. The Tolman length physically describes the difference between the surface of tension R_s and the equimolar dividing surface R_e,

$$\delta_T \equiv \lim_{R_s \to \infty} (R_e - R_s). \tag{2.56}$$

The surface of tension is where interfacial tension is defined as physically occurring within a Gibbsian framework between the particle and the surrounding phase or where the standard Laplace equation is valid ($R_s = 2\sigma(R_s)/\Delta p$). The equimolar dividing surface is where $4\int_0^\infty [\rho(R) - \rho_v] R^2 \mathrm{d}R = (\rho_l - \rho_v)\frac{4\pi}{3} R_e^3$. Here ρ_l and ρ_v are the bulk liquid and vapor densities for the case of a liquid-vapor interface. Holten (2005) [62] derived, using the

capillary approximation and assuming Tolman's approximations, that R_s is $2\delta_T$ from the capillary radius R_{cl}, or radius at which bulk properties are valid and which is commonly used in classical nucleation theory, and R_e is directly between R_{cl} and R_s. One means of implementing the Tolman equation is by utilizing its Taylor series expansion about $r = \infty$ truncated to two terms,

$$\sigma(r) = \sigma_{\text{bulk}}(1 - 2\delta_T/r). \tag{2.57}$$

This truncated series form will be implemented within the current nucleation model along with its associated effects upon coarsening, as was previously alluded to in Section 2.3.3.

2.4 Solid-Liquid Phase Interaction

The solid and liquid phases of the precipitating system are interdependent and need to be coupled. Solid-liquid coupling can be accomplished through the use of an extent of reaction variable χ. This reaction-extent variable represents the normalized total amount of solid material precipitated out of solution. The amount of solid material present in the system at a position in time can be calculated by summing the third moments of all polymorphs, representing total volume, multiplied by their respective densities,

$$c_j v_j = \frac{4\pi c_j}{3} \int_{r_1}^{\infty} r^3 \eta_j \, \mathrm{d}r = \frac{4\pi c_j}{3} m_{3,j}, \tag{2.58}$$

and normalizing by the total amount of material initially available in the system,

$$\chi = \frac{4\pi}{3} \sum_j \frac{c_j m_{3,j}}{n_{\text{CaCO}_3,\text{max}}}. \tag{2.59}$$

Here c_j represents polymorph j's molar density, v_j is the polymorph's total volume, and $n_{\text{CaCO}_3,\text{max}}$ is the maximum amount of material that can be precipitated out of the system, which can be determined from the stoichiometrically limiting ion. Chemistry tables relating chemical activities to reaction-extents can be created *a priori*. Thus, as material precipitates into solid forms at each time step, the extent of reaction can be calculated and noted to be proportional to the sum of the rates of change of the polymorph's third moments

$$\frac{\mathrm{d}\chi}{\mathrm{d}t} \propto \sum_j c_j \frac{\mathrm{d}m_{3,j}}{\mathrm{d}t}. \tag{2.60}$$

After each time step the aqueous-phase's chemical activities are updated by interpolation within the chemistry tables depending upon the current extent of reaction.

2.5 Conclusions

To capture the evolution of a highly supersaturated aqueous system as it evolves to thermodynamic equilibrium, phases and timescales had to be coupled. Equilibrium chemistry generated by an ionic chemistry package was tabulated over the range of reaction extents. Population balance equations were derived to capture the coupling of solid-phase processes nucleation, growth, aggregation, and dissolution. The coupling of aqueous and solid-phases through the reaction-extent allowed thermodynamics effects such as Ostwald ripening to be captured. A simplistic mixing model was also implemented to roughly emulate mixing effects. Lastly, thermodynamic relations were derived to capture temperature and composition effects upon interfacial tension. The aggregate of all these modeling pieces embodies a modeling framework for capturing the evolution of $CaCO_3$ systems of interest and will function at spacial and temporal scales used within the CFD applications.

CHAPTER 3

PRECIPITATION FRAMEWORK RESULTS AND ANALYSIS

3.1 Introduction

The validation metric used to test the precipitation framework created within Chapter 2 was a dataset from Ogino et al. [106]. The experiments described within Ogino et al. were reported to have been carried out at fixed temperatures and were chiefly initiated at a single set of initial conditions. With this data for systems with fixed temperature and initial composition, the precipitation framework was able to produce results that were directly comparable to those reported. Experimental data were presented within Ogino et al. primarily in the form of volumetric polymorphic abundance and ion-activity product (IAP) traces. IAP traces can be extracted from the equilibrium-chemistry tables as the product of $Ca_2{}^{+}$ and $CO_3{}^{2-}$ activities. Each polymorphic form's third moment can be utilized to track that form's total volume, which in comparison to the total solid volume, or sum of the four forms, describes polymorphic abundance traces.

Results of this validation step will be analyzed in a step-wise fashion similar to that utilized throughout the framework's development. Many potential configurations of the physics submodels comprising the framework are possible. Analyzing the effect of a subset of the potential configurations while in the process of choosing a final configuration for a more thorough validation will be presented first. The interfacial tension parameters were known to have large amounts of uncertainty and the system's evolution was highly sensitive to their values. Once a final framework configuration was chosen, the interfacial tension values were optimized and framework outputs compared with the data from Ogino et al. [106]. Although only two types of data were utilized in this framework validation, the precipitation framework has the ability to produce many additional forms of data and

---

Adapted with permission from B. Schroeder, D. Harris, S. Smith, and D. Lignell, *Theoretical framework for multiple-polymorph particle precipitation in highly supersaturated systems*, Cryst. Growth Des., 14 (2014), pp. 1756-1770.

some of these capabilities will be analyzed to extract additional insights. Before framework analysis is presented, a description of a hypothesized error in the experimental data as well as how it is corrected are outlined.

3.2 Corrections to Literature Data

During the analysis of the experimental data presented within Ogino et al. [106], it was hypothesized that an error was present within the BASIC code utilized by the experimentalists when calculating the published results. Although it was reported that the BASIC code was based upon the chemistry model developed by Plummer and Busenberg [109], it can be seen in Figure 3.1 that a reimplemented version of Plummer and Busenberg's model did not match the reported data of Ogino et al. By changing the model such that the direction of calcium carbonate complex, $CaCO_3°$, formation was effectively reversed, Ogino et al. data were matched. The implementation of Plummer and Busenburg's model was then checked against calcium concentrations measured by Gebauer et al. [42] in a similar experimental setup.

Following the assumption that the hypothesized error in the Ogino et al. results was correct, accounting for such an error is simple and allows the data to be utilized for framework validation. Figure 3.2 demonstrates how the experimental data can be corrected through correlation to the original model and the chemistry produced when using

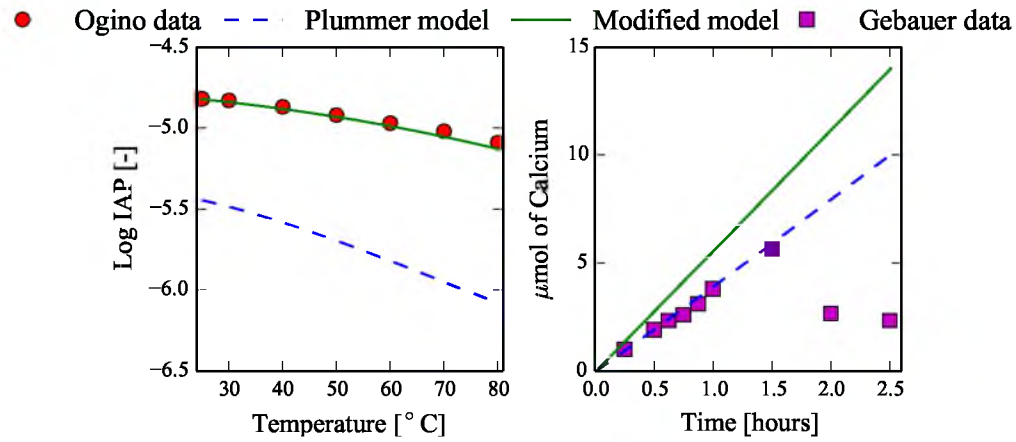

Figure 3.1: Demonstration of hypothesis that a model with a reversed equilibrium expression was likely utilized within results reported by Ogino et al. [106]. Dashed lines represent the original $CaCO_3$ aqueous ionic-chemistry model developed by Plummer and Busenberg [109] and solid lines represent modifying that model such that the $CaCO_3°$ equilibrium expression is reversed. The left plot includes IAP data from Ogino et al. (circles) over a temperature range from 25 to 80°C, while the right plot contains calcium concentration data over time as reported by Gebauer et al. [42] (squares).

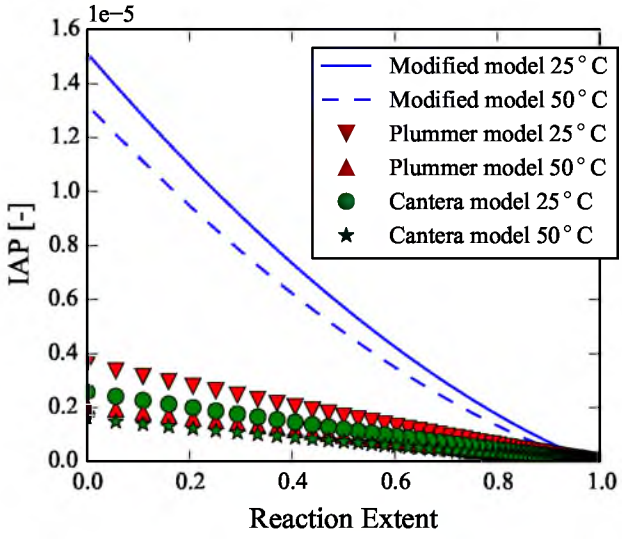

Figure 3.2: IAP traces over the full range of reaction extent, starting with the initial conditions used by Ogino et al. [106] for systems at 25°C and 50°C, are compared for three ionic-equilibrium chemistry models. The model presented in Plummer and Busenberg [109] is shown as downward and upward pointing triangles, Plummer and Busenberg's model with the $CaCO_3°$ equilibrium expression reversed is shown as solid and dotted lines, and the equivalent values calculated using Goodwin et al. [50] are shown as circles and stars, all for 25°C and 50°C, respectively.

the Cantera package [50] as a function of reaction extent. Such correlation was utilized throughout all portions of the framework validation. It should also be noted that this hypothesized error had no bearing upon the timescale and particle population statistics reported by Ogino et al.

3.3 Configuration Analysis

In order to evaluate the performance of the precipitation framework, a configuration of physical submodels had to be chosen. It was desired that the framework remain fairly simple in order to avoid overfitting the limited amount of data available, while still being able to capture the system's dynamic trends across multiple time-scales. Many potential submodels were presented within Chapter 3. Beyond submodel selection, numerical implementation issues also had to be addressed.

When utilizing moment methods to solve the PSDs, the number of quadrature nodes utilized must be selected. Ideally, the more nodes evolved the better, but additional nodes add computational cost, thus creating a balancing problem. Initial simulations were run with two and three nodes. While producing similar results, there was enough discrepancy

to justify exploring the use of more nodes. Using four or five nodes (eight or ten moments) for each PSD did not cause notable differences, when compared to the results produced using three nodes. Thus, it was decided that evolving three nodes was the balancing point between accuracy and computational speed.

The mixing-time, or time until complete mixing of the two reactant streams, also needed to be selected. Once the optimization process located the vicinity of likely parameter values, a range of mixing times were surveyed. It was found that a mixing-time that allowed for the majority of mixing to be concluded by 1.5 seconds and complete mixing prior to 5 seconds was optimal. The actual value of the mixing-time model-parameter utilized in the three-environment mixing-model Eq. (2.30) was 0.75 seconds. This mixing time allowed the optimization scheme to locate a set of parameters that corresponded to an acceptable amount of discrepancy from the experimental data. Additionally it was judged to be a physically realistic time-frame for the system's specified mixing geometry. Supplemental data on the PSD and number density at short time-scales would allow for a better understanding of the mixing event and potentially justify the implementation of a more detailed mixing model. The other time-scale, the time-lag parameter, utilized for transient nucleation [69] was found to have no notable effect for the system conditions of interest, but might for systems with lower concentrations.

Also within the nucleation model, diffusion-limited and interface-limited reaction-rate coefficients were compared. The two reaction-rate coefficients were found to produce similar results and corresponded to approximately 0.01% and 2.0% differences in the interfacial-tension reference points and excess entropy values found during optimization. Such similar behaviors were also reported by Lindenberg and Mazzotti [83]. Ultimately, the diffusion-limited variant was selected for the final framework configuration to maintain consistency with the diffusion-limited growth mechanisms utilized when the majority of nucleation occurred.

The initial framework configuration considered consisted of homogeneous nucleation, aggregation, mixing, and diffusion-limited kinetics for growth and dissolution for all $CaCO_3$ forms except ACC's dissolution, which used a kinetically limited rate in the form of Eq. (2.20) with a rate constant $9 \times 10^{-9}\, m\, s^{-1}$. This configuration was compared with and parameters fitted to experimental data for 25°C, as is shown in Figure 3.3. Although this initial framework configuration captured the slopes, time-scales, and general trends reasonably, its primary purpose was to act as a basis of comparison throughout framework development.

Even though the basic framework configuration performed well in matching the exper-

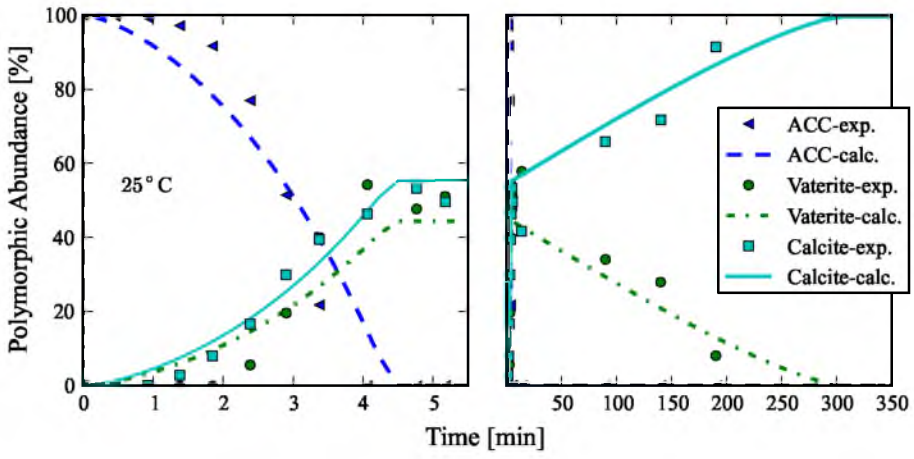

Figure 3.3: Traces of polymorphic abundance over time at 25°C calculated by the precipitation framework (lines) compared with experimental data (symbols) extracted from Ogino et al. [106]. The interfacial tension values used within the precipitation framework were optimized to fit the experimental data at this temperature with greater emphasis upon the shorter time-scale data. ACC data are shown as the dashed lines and triangles, vaterite data are shown as the circles and dash-dot lines, and calcite data are shown as the solid lines and the squares for the calculated traces and experimental data, respectively. At these conditions aragonite is not present in any significance.

imental data, there are still inconsistencies that can be used as points of comparison and analysis while refining the configuration. In order to capture the general polymorphic abundance trends and time-scales, the initial period of ACC dominance has been cut short. The time-lag between when calcite becomes prolific and when vaterite also reaches similar quantities has not been captured. In order to capture the short time-scale trends, simultaneous vaterite and calcite increasing trends were required. Lastly, this configuration did not allow vaterite and calcite to switch polymorphic abundance dominance roles twice, once in the short time-scale and once in the long time-scale. That being said, all of these criticisms should not be given too much weight considering the scale of potential error in the experimental data was reported to be up to 10%, potentially invalidating most of these issues. Throughout the exploration of framework configurations, the two most fundamentally significant options explored dealt with the selection of nucleation and growth mechanisms.

Within Figure 3.4 simultaneous heterogeneous and homogeneous nucleation has been implemented. Both forms of nucleation are likely to have occurred within the system of interest, but their relative levels of importance are unknown. The simultaneous implementation does a notably better job at capturing most of the perceived shortcomings of the basic framework configuration. ACC is able to maintain near 100% polymorphic abundance for

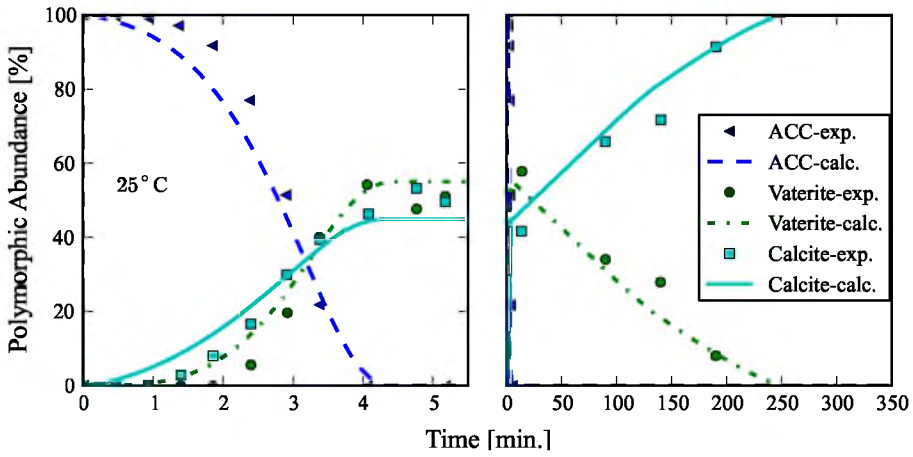

Figure 3.4: Traces of polymorphic abundance over time at 25°C calculated by the precipitation framework (lines) compared with experimental data (symbols) extracted from Ogino et al. [106]. Combined heterogeneous and homogenous nucleation with parameters fitted to best capture the experimental data were used with diffusion-limited growth and dissolution mechanisms for all $CaCO_3$ forms. ACC is shown as dashed lines and triangles, vaterite is shown as dash-dot lines and circles, and calcite is shown as solid lines and squares, for the simulated and experimental data, respectively.

a longer period of time. A period of lag between calcite's rise in polymorphic abundance and vaterite's exists, albeit occurring a minute premature. Even the short- and long-term intersections of polymorphic abundance for vaterite and calcite are captured. Although such results do appear promising, the heterogeneous nucleation mechanism introduces additional unknown parameters, which along with the interfacial tension values were optimized for fitting the experimental data. Optimizing too many parameters can add too many degrees of freedom to the framework for the results to be considered conclusive. Thus, to avoid such overfitting only homogeneous nucleation will be utilized in the final configuration.

A wide variety of permutations of growth and dissolution mechanisms exist for the framework including the possibility of transitioning between different mechanisms. Figure 3.5 demonstrates how a configuration that includes a transition in growth mechanisms is able to capture dynamics absent in the basic configuration. For Figure 3.5 the only change from the basic configuration was that vaterite's growth mechanism shifts from diffusion-limited to screw-dislocation based at a supersaturation value of 47 and the interfacial tension values were refitted.

Transitioning vaterite's growth mechanism allowed for the delay in vaterite reaching significant polymorphic abundance and overtaking calcite around 3 minutes to be captured. This is caused by a change in the fitted interfacial tension values of vaterite relative to

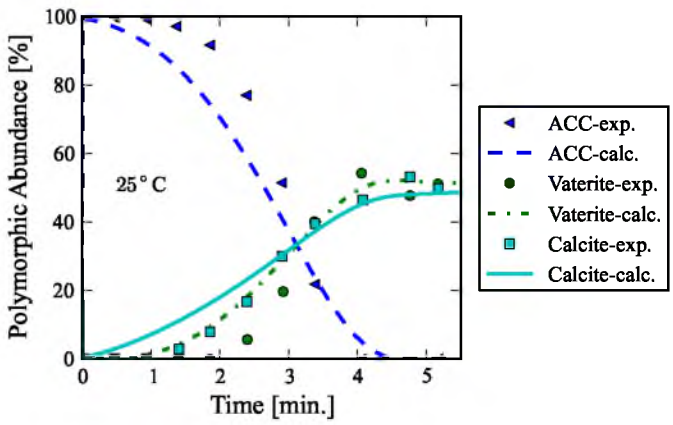

Figure 3.5: Traces of polymorphic abundance over time at 25°C calculated by the precipitation framework (lines) compared with experimental data (symbols) extracted from Ogino et al. [106]. Vaterite's growth mechanism switched from diffusion-limited to a screw-dislocation mechanism when its respective supersaturation became less than 47. All other growth and dissolution mechanisms were diffusion-limited and interfacial tension values were fitted to best capture the experimental data. ACC is shown as dashed lines and triangles, vaterite is shown as dash-dot lines and circles, and calcite is shown as solid lines and squares, for the simulated and experimental data, respectively.

the other polymorphs. Realistically, the implemented change does not represent the true physical transition occurring because there are multiple growth mechanisms that should be transitioned through after diffusion-limited growth and prior to a screw-dislocation mechanism becoming dominant [27]. At the same time, inclusion of the full range of growth mechanisms was not desired in order to keep the framework relatively simple. The chosen transitioning point was arbitrary and will make extrapolating this configuration to other temperatures likely to fail. Another issue with this dual growth-mechanism configuration is that fitting interfacial tension values to allow for good short time-scale trend-matching causes the long time-scale trends to extensively overshoot the times of the experimental data. If multiple growth mechanisms were ever utilized, it appears likely that more than two would be necessary as well as transitions in the dissolution mechanisms.

Recently there has been a great deal of research on calcite's growth and dissolution rates and many complex mechanisms have been suggested [82, 132, 43, 18, 111, 119]. Wolthers et al. [145] presented a calcite growth mechanism that accounted for ionic ratio and pH dependence. This mechanism was found to perform well when transitioned to from the diffusion-limited mechanism at a supersaturation value around 1.5. In order to incorporate temperature dependence within calcite's dissolution, the Plummer et al. [110] mechanism can be utilized.

3.4 Final Configuration

After considering many possible model configurations, one was ultimately selected as the best representation for the system of interest. This final configuration included mixing, homogenous nucleation (without time lag), aggregation, diffusion-limited growth except for calcite at supersaturations below 1.5 where the Wolthers et al. [145] mechanism is used, kinetically-limited dissolution for ACC, diffusion-limited dissolution for vaterite and aragonite, and the Plummer et al. [110] dissolution mechanism for calcite. A Nelder-Mead Simplex optimization was used to locate the slope and intercept parameters of the interfacial tension model Eq. (2.54) for each polymorph. The error kernel that the optimization scheme explored was based upon a weighted sum of squared-error between the framework outputs and the experimental polymorphic abundance and IAP data across all temperatures provided by Ogino et al. The optimized interfacial tension parameters are tabulated in Table 3.1 and Figure 3.6 depicts how the framework performs with optimized parameter values over the full range of temperatures. The comparison of the framework output with the experimental data points provides a visual of a portion of the error kernel minimized.

Within Figure 3.6 there are clearly areas of varying success in the framework's ability to capture the experiment data trends. Short time-scales and temperatures where data were available across the full time-spectrum are captured more satisfactorily. The worst performance is for the long time-scales at 80°C where the framework's predicted time-scales far exceeded the experimental data. General trends for all polymorphic forms across the full temperature range were captured, reflecting the framework's ability to generally account for temperature effects. Short-term data matching was likely more successful due to being more directly related to nucleation, which is most strongly affected by the optimized interfacial tension parameters.

A great metric for the framework's ability to capture temperature related trends can be found in Figure 3.7, where the metastable polymorphic abundances from the onset of the metastable period are plotted. The onset of the metastable period was interpreted as

Table 3.1: Interfacial tension model parameters (for Eq. (2.54)) from optimization using the Nelder-Mead Simplex method with weighed residual.

Polymorph	$\sigma_{25^\circ C,\mu^\circ}$ [mJ m^{-2}]	$-S_a^E$ [mJ m^{-2} K^{-1}]
ACC	-12.6	2.47×10^{-2}
Vaterite	64.6	3.57×10^{-2}
Aragonite	96.3	-3.16×10^{-1}
Calcite	8.94	1.06×10^{-1}

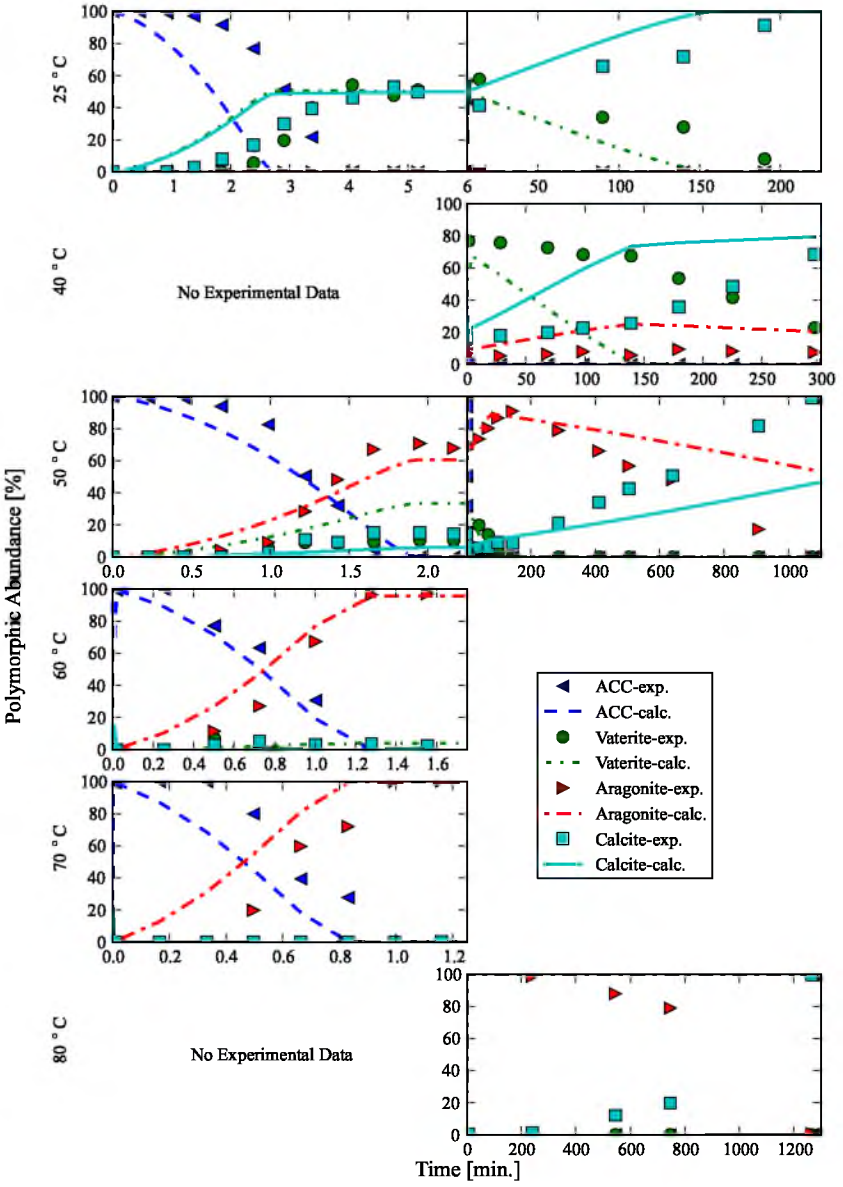

Figure 3.6: Traces of polymorphic abundance over time calculated by the precipitation framework (lines) compared with full temperature range of polymorphic abundance experimental data (symbols) extracted from Ogino et al. [106]. The final framework configuration as described in Section 3.4 was utilized with optimized interfacial tension values. ACC is shown as dashed lines and left-pointing triangles, vaterite is shown as shorter dash-dot lines and circles, aragonite is shown as the right-pointing triangles and longer dash-dot lines, and calcite is shown as solid lines and squares, for the simulated and experimental data, respectively. For 80°C aragonite's IAP trace is approximately 100% for the entire time range shown.

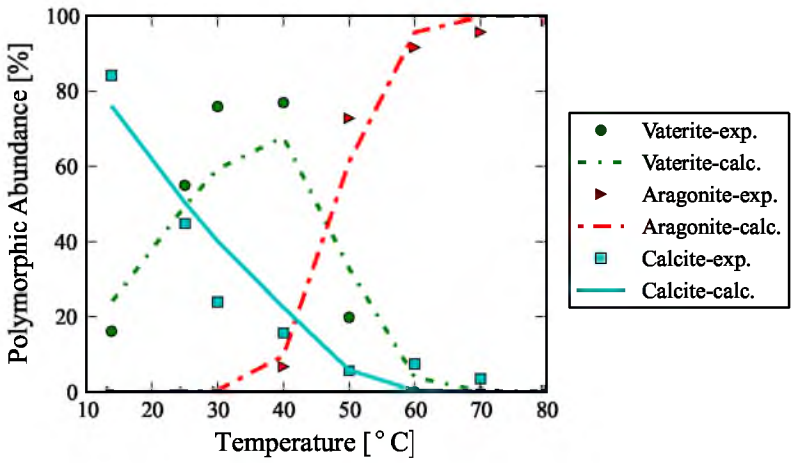

Figure 3.7: Polymorphic abundances calculated at the metastable stage by the precipitation framework (lines) compared with experimental data (symbols) extracted from Ogino et al. [106] over a range of temperatures. The time at which the metastable stage is reached was defined as the when the system's IAP levels out following the majority of ACC dissolving. Vaterite is shown as the dash-dot line and circles, aragonite is shown as the triangles and longer dash-dot line, and calcite is shown as solid lines and squares, for the simulated and experimental data, respectively. The times at which data were extracted from the precipitation framework was set to those tabulated within Ogino et al., except at 13°C, where a longer time was used due to significant ACC still being present at the tabulated time.

occurring once the majority of ACC has dissolved and the IAP traces reached their second flat region. Within Figure 3.8 the transition to the metastable period can be observed to occur near 6 and 3 minutes for the 25°C and 50°C experimental data, respectively. Metastable onset was tabulated across temperature within Ogino et al. and those times were primarily used to define this same point for the framework's output. Comparing the framework's predictions with the experimental data demonstrates that the the interfacial tension model has satisfactorily captured temperature dependence. The metastable period is reached faster at higher temperatures and it is there that the framework best resembles the experimental data. At lower temperatures the trends for each polymorph are captured, but are dampened in amplitude. This generally shows that the framework's response is less sensitive to temperature changes than was experimentally observed. Considering the approximations and simplifications known to be utilized within the physical submodels, the framework's performance at capturing the metastable stage was deemed reasonable.

To observe the framework's performance on all time-scales, the IAP traces in Figure 3.8 can be analyzed. Figure 3.8 compares the outputs calculated by the framework with the raw experimental data extracted from Ogino et al. [106], as well as that experimental data once

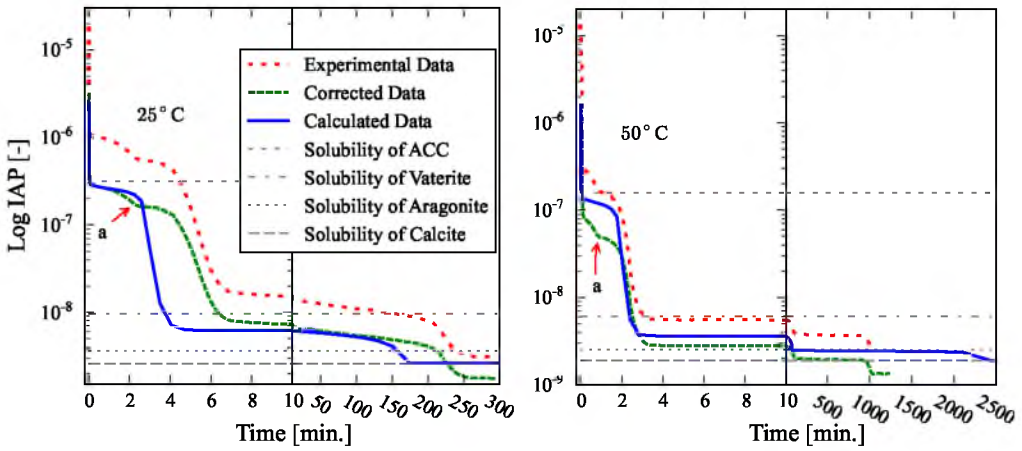

(a) (b)

Figure 3.8: IAP traces over time at 25°C (a) and 50°C (b) compared with experimental data extracted from Ogino et al. [106]. The experimental data as originally extracted from Ogino et al. are shown as wide-gapped dotted lines, the extracted data when corrected as depicted within Figure 3.2 were shown as narrow-gapped dotted lines, and the data calculated by the precipitation framework were shown as solid lines. The solubilities of the four forms of $CaCO_3$ are included for comparison with metastable periods of the particle evolution and are shown as vertical lines where ACC's are short dashed lines, vaterite's are dash-dotted lines, aragonite's are dotted lines, and calcites are long dashed lines. Arrows direct attention to inflection points, labeled 'a,' that are caused by an unknown particle phenomena and not captured by the particle framework.

the hypothesized error was corrected for using the correlation with the Pitzer equation based chemistry equivalent [50]. The first discrepancy between the original experimental data and those calculated by the framework is the value at which the IAP initiates. This was a clue that eventually lead to the hypothesis of an error in the code utilized to create the reported experimental data. Another issue with the extracted experimental data is that the IAP remains above ACC's solubility line until ACC dissolves. For systems containing multiple polymorphs it would be expected that the IAP values will remain below the solubility of the polymorph that is currently least stable. Once the correlation correction is applied, the experimental data satisfy this requirement.

Following ACC's dissolution the IAP stabilizes into a metastable period at a value below the next most thermodynamically stable metastable polymorph vaterite. As was mentioned previously, the point at which the IAP reaches this second point of stability has been termed the metastable point, which was compared within Figure 3.7. This process continues on until the thermodynamically favored polymorph calcite is reached. At 50°C a transition through aragonite occurs, but at 25°C this polymorph is insufficiently stable to be present

in significant enough quantities to impact the chemistry.

There appears to be an issue with the corrected experimental data during the longer time-scale data when aragonite and calcite become dominant. The corrected IAP traces level off near calcite's solubility when aragonite is still present in significant quantities within the system and eventually moves below calcite's solubility. It would be expected that the final IAP would asymptote at calcite's solubility as the system approaches equilibrium due to calcite being the only polymorph remaining. Many potential explanations exist for why this is observed, one being possible fouling of the probe.

It appears that the current framework configuration was able to capture all of the general IAP and polymorphic abundance behaviors reasonably well considering the use of basic models and many assumptions. Of particular note is that the short-term time-scale trends were captured across a wide range of temperature while primarily using diffusion-limited growth and dissolution mechanisms (as can be seen within Figure 3.6). This is noteworthy considering it was known that other growth mechanisms occur within $CaCO_3$ systems, but it appears that with high initial supersaturation, this approximation was acceptable. Calcite did have additional kinetic mechanisms enacted for its growth and dissolution, but due to the conditions in which they were utilized they should have only minor effects. For greater consistency with the long-term data, additional growth and dissolution mechanisms could be implemented as a future endeavor.

In both experimental IAP datasets extracted from Ogino et al. [106], there are two time-periods characterized on the basis of uniform time sampling: with order of tens of seconds for the short time-scales and tens of minutes for the long time-scales. While this adequately depicts the system's general trends, it does little to aid in characterizing nucleation and the initial mixing event. Within Ogino et al., perfect mixing was assumed, so it is not surprising that this time frame was not better characterized. In the IAP traces for 25°C and 50°C there are large IAP drops initially, but they appear to occur at time zero. If greater time resolution were available the nucleation and mixing models could have been more thoroughly validated, effectively removing a significant amount of uncertainty from the framework.

Another feature of concern with the IAP traces are the unexplained inflection points, located approximately in the middle of the ACC's period of stability, which are labeled in Figure 3.8 with an 'a.' The current framework configuration was unable to recreate these inflection points, but many potential physical causes were hypothesized. A secondary form of ACC, as was discussed in Cartwright et al. [21] and Radha et al. [112], could exist within

the system with a lower solubility than the variety currently assumed. Further support of the potential existence of a second form of ACC can be found in position of the inflection point. The infection point occurs under ACC's known solubility for both 25°C and 50°C, but well above that of vaterite. All other inflections in the IAP traces correspond to the system reaching a new metastable phase, so it seems reasonable that this point would also.

An additional noncaptured physical phenomena potentially related to that observed in the IAP traces can be found in the polymorphic abundance trances shown in Figure 3.6. Until near 2.0 minutes at 25°C, 0.75 minutes at 50°C, and 0.3 minutes at 60°C and 70°C, ACC encompasses nearly 100% of the polymorphic abundance. This phenomena can be viewed either as a longer period of ACC stability or a delay in the nucleation/growth of the other polymorphs. Either way the current framework configuration does not capture this behavior. The possibility of multiple ACC forms, as was suggested for the noncaptured IAP inflections, could be the cause of this considering that the second form of ACC would be expected to have lower solubility than the variety of ACC currently considered and thus remain stable longer. Transient nucleation was initially considered as a potential source of the delay, but induction-times were found to not be within the order of magnitude necessary to correlate with the observed phenomena. Alternative nucleation mechanisms for ACC and heterogeneous nucleation were also entertained as potential causes.

While the cause of the noncaptured physical phenomena could not be definitively proven, a temperature functionality was clearly observed. The event occurred on longer time-scales at lower temperatures. Where the IAP inflection occurred around 2 minutes at 25°C, it occurred in half that time at 50°C. The scaling of this effect is in line with the temperature effects upon the transition from ACC to the next most stable polymorph.

At both 25°C and 50°C there appear to be two points at which vaterite and calcite switch prevalence of relative polymorphic abundance. In Figure 3.6 the transitions for 25°C can be seen to occur around 3.5 and then between the first two long time-scale experimental data points. The short time-scale transition at 50°C cannot be directly observed in the available data, but can be inferred between comparing the short and long time-scale polymorphic abundance traces along with the metastable data in Figure 3.7. At 50°C the second transition occurs around 600 minutes. The long time-scale transitions can be attributed to calcite being the thermodynamically most stable polymorph. What causes the short time-scale transitions is less clear, but changes in the polymorphs growth mechanisms as the system's conditions change seems likely, as was previously noted in Section 3.3.

3.5 Additional Output Analysis

Although validation and calibration of the framework based upon comparison with the experimental data of Ogino et al. [106] was limited to the data forms included within that article, the framework developed allowed for many additional forms of data output. Evolving moments of the particle distribution allows for means, variances, and other statistical qualities of the particle distribution to be characterized across time. This provides a wealth of potential data forms for comparison with other experimental data sources. Examples of such potential statistical characteristics include the particle number density, average particle size, and variance in particle size as shown in Figure 3.9-3.10. The simultaneous tracking of the solid-phase PSD and aqueous-phase ionic-equilibrium chemistry allows for correlations between the events occurring in each phase to be extracted and analyzed, as will be done with the supersaturation traces in Figure 3.10.

An initial period of nucleation and mixing until around 10^{-7} minutes can be observed in both the plots of the particle number density and average radial size shown in Figure 3.9. After this period of nucleation, the number of particles remains stagnant while the particle size increases, indicating that growth and mixing are occurring. ACC is a slight exception to this trend until 10^{-4} minutes while aggregation is occurring. The fact that only ACC experiences significant aggregation within this time period should not be surprising considering that aggregation is a function of number density and ACC has approximately five orders of magnitude more particles in solution following nucleation than any other $CaCO_3$ form.

These same temporal periods can be further analyzed through the particle distribution's normalized standard deviations and supersaturation traces shown in Figure 3.10. During the nucleation and mixing-limited times, the radial distributions' variances increase corresponding to the fact the size at which particles are nucleating into the distributions, the critical size, is shifting. Although the variance is increasing, these are still relatively narrow distributions. Once growth becomes the dominant phase-transformation mechanism, the distribution's variances slope downwards. The downward slope is caused by diffusion-limited growth being size dependent, causing smaller particles to grow faster than larger particles. Such size dependent behavior is effectively collapsing the PSDs. Again, ACC is an exception with aggregation causing a continued increase in spread of the radial distribution. Vaterite and calcite also show signs of minor aggregation between 10^{-4} - 10^{-2} minutes within the average radial size plot.

Within the supersaturation traces the onset of growth corresponds to the first downward

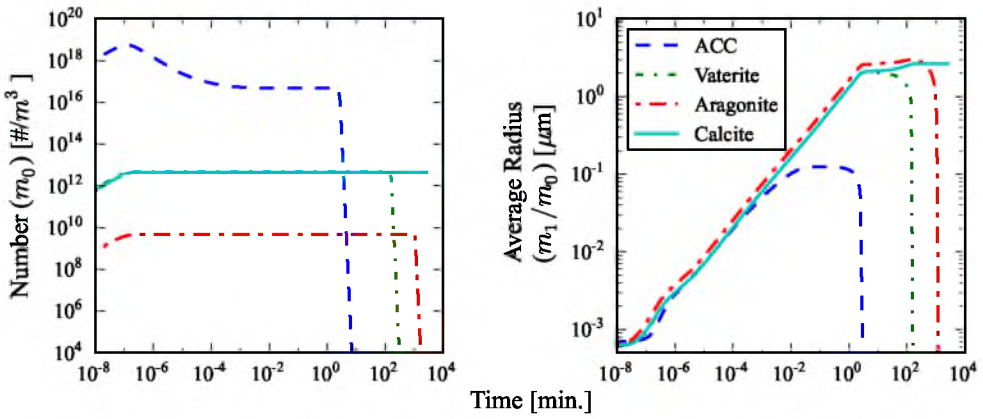

Figure 3.9: Traces of the particle number density or 0^{th} moment, m_0, over time in the left plot and traces of the average radial size, m_1/m_0, over time in the right plot for all four forms of $CaCO_3$. ACC is shown as the dashed lines, vaterite is shown as the shorter dash-dotted lines, aragonite is shown as the longer dash-dotted lines, and calcite is shown as the solid lines.

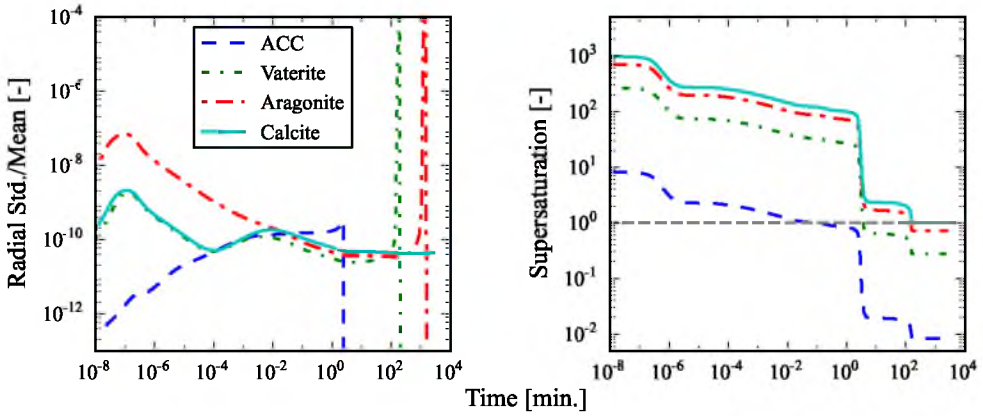

Figure 3.10: Traces of the polymorphic PSD standard deviation normalized by their respective means over time are shown in the left plot and the respective supersaturation traces over time are shown in the right plot. A horizontal dashed line is shown in the supersaturation plot to provide reference for where a supersaturation of one is located. ACC is shown as the dashed lines, vaterite is shown as the shorter dash-dotted lines, aragonite is shown as the longer dash-dotted lines, and calcite is shown as the solid lines.

shift. No shift of similar proportion was caused by the entire nucleation process, indicating that growth has greater control over the system's supersaturation and is the dominant means of phase change occurring within the system. The effect of mixing upon particle growth can be noted by the bump in the supersaturation traces between 10^{-6} - 10^{-2} minutes. It is during this time that the mixing of the reactant streams is limiting the speed at which growth is occurring.

Shifting focus onto only ACC's traces, once its supersaturation begins to level out near unity there is no longer a driving force for growth. This effects smaller particles before larger particles due to the $\bar{S}$ term used within the growth mechanisms. As the smaller ACC particles begin to dissolve, the radial variance experiences a slightly positive slope. Then as larger portions of the PSD begin to experience dissolution, Ostwald ripening causes a spike in the radial variance. The continued growth of the other polymorphs eventually causes dissolution to become energetically favorable for all sizes of ACC and the radial variance drops off as ACC effectively disappears from the system. Within the average radius plot, positive inflection of the more stable polymorph's slopes corresponds with ACC's dissolution. The coupling of the polymorphic forms' PSDs is able to be captured in this manner due to the coupling of all solid forms with the aqueous-phase ionic-equilibrium chemistry table. Without such coupling ACC's dissolution would not be forced by the growth of the other polymorphic forms, reflecting their favorable thermodynamic stability.

Trends similar to ACC's dissolution can be seen for vaterite and aragonite. Distribution coarsening causes more drastic spikes in these metastable polymorphs' radial variances. The larger amplitude of the spikes are due to the larger mean size of the polymorphs prior to dissolution which allowed for a wider distribution to be formed as the side of the distribution corresponding to smaller particles elongated. Due to aragonite's significantly lower number density (for the depicted conditions), there is no drop in the supersaturation traces corresponding to its dissolution. Once all metastable polymorphs have dissolved away, the thermodynamically favored calcite remains with a relative supersaturation around unity. Little will change in the system as calcite slowly experiences Ostwald ripening until equilibrium is eventually reached, but this is well outside the time frame of interest.

Through comparing the number density and the average radial size traces, the sequence of dissolution events can be confirmed. The average radial size and variance both drop significantly prior to the number density's decline for each of the metastable polymorphs. It should be stressed that the linear slopes corresponding to polymorph dissolution shown in the number density plot are due to the nonphysical death kernel implemented, which

was designed to have a linear slope in log-log space. Although this is not based directly upon a physical mechanism, it is physical in that the particles have already significantly shrunk and should be reentering the aqueous-phase. As previously noted, particle death only occurs once it has been ensured that it will have a negligible effect upon the chemistry and remaining polymorphs.

While tabulated particle size characteristics were not reported by Ogino et al., figures containing electron micrograph photos were shown for a few system conditions. The Ogino et al. Figure 5-h shows vaterite and calcite particles 50 minutes into the process at 25°C. The depicted particles had radii around two microns, which is in reasonable agreement with the calculated averages shown in Figure 3.9. Ogino et al. also shows in Figure 5-e ACC particles at 7 minutes into the process for the same conditions with radii ranging from 0.5 to 2.0 μm. While this range does not overlap with the calculated average radii of $\sim 0.1\,\mu m$, it is within an order of magnitude which is considered satisfactory given the large amount of uncertainty within the system's submodels.

3.6 Parameter Uncertainty

Framework results presented thus far utilized calibration parameters fitted with a weighted-sum squared-error minimization technique. While the framework's results appeared reasonable in comparison with experimental data, the plots shown did not demonstrate the framework's sensitivity to the calibration parameters. To clearly see this an exploratory sensitivity study was conducted. For this sensitivity study a centered Latin hypercube DOE was utilized to select ten parameter sets. The bounds of the parameter space explored were plus or minus one percent around the optimally fit interfacial tension parameters, shown in Table 3.1, and plus or minus ten percent around the chosen mixing time, 0.75 seconds. The IAP and polymorphic abundance traces at 25°C for the ten calibration parameters sets are shown in Figure 3.11-3.12.

The effect of the specified parameter variance when propagated through the framework was clearly evident in the IAP traces. While in short time-scales minimal differences can be noted, by the first significant dissolution event the times at which the IAP drops occur have spread. The IAP traces then restabilize at varying IAP values, that do not directly correlate with the previously noted variance in IAP drop time. During the second major dissolution event there is also variety of time-scales over which this is observed and the temporal variance within the dissolution event now corresponds to time on the order of tens of seconds. The IAP traces do all still maintain similar trends and equilibrate to the same

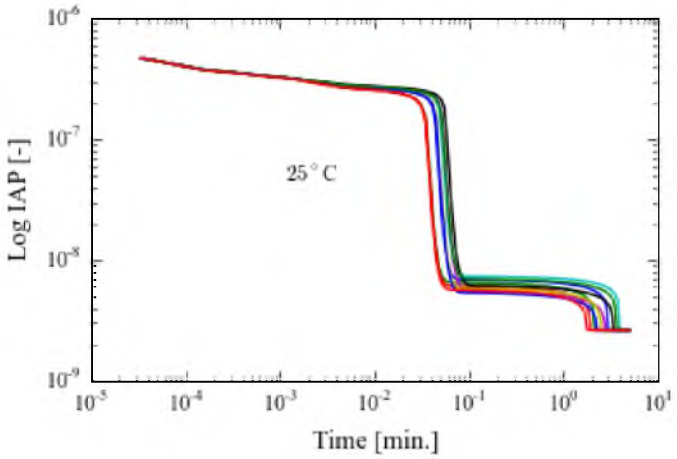

Figure 3.11: Ten IAP traces over time at 25°C that were created by varying the interfacial tension values and mixing times. The interfacial tension and mixing time values were varied using ten points from a centered Latin hypercube design were the parameter space varied plus or minus one and ten percent, respectively, around values previously found to optimally fit experimental data from [106].

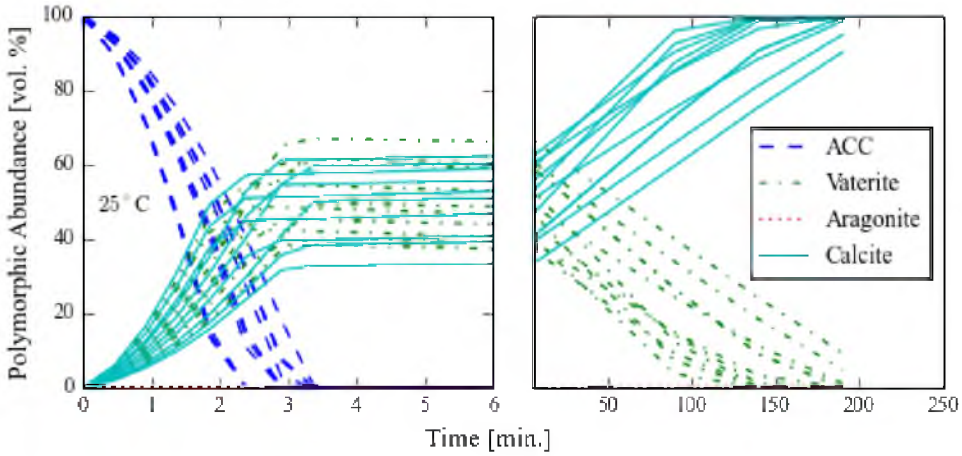

Figure 3.12: Ten instance of polymorphic abundances traces over time at 25°C that were created by varying the interfacial tension values and mixing times. The interfacial tension and mixing time values were varied using ten points from a centered Latin hypercube design were the parameter space varied plus or minus one and ten percent, respectively, around values previously found to optimally fit experimental data from [106]. ACC traces are the dashed lines, vaterite are the dash-dotted lines, aragonite are the dotted lines, and calcite are the solid lines.

final value.

Propagation of the parametric uncertainty also shows significant effects upon the polymorphic abundance traces. For short time-scale polymorphic abundance traces, the variance in the calibration parameter causes the time to complete ACC dissolution to spread over a minute and likewise the time to metastability for calcite and vaterite. Vaterite and calcite's metastable polymorphic abundances vary over 30% and sometimes swap relative prevalence due to the uncertainty propagation. These short time-scale effects translate into variance in the long-time scales of over 100 minutes. Again, while time-scales and polymorphic abundance values are effected by the parametric uncertainty, the general shapes and trends are not changed.

Clearly the framework and the results by which it is being validated are largely sensitive to the calibration parameters. At least near the optimized parameter values, the effect of varying the parameters appears to produce smooth changes and does not drastically alter the system trends. Given that this sensitivity study only explored a small region of the true uncertainty contained within the parameter, especially those describing interfacial tension, this system clearly would benefit from a thorough uncertainty quantification. A study where interfacial tension effects could be better isolated from other system physics, such as a nucleation study, could allow better calibration of the parameters.

3.7 Conclusions

The particle precipitation framework developed within Chapter 2 was validated against experimental data extracted from Ogino et al. [106]. First these data were studied and a hypothesis was created that an error was present in the presented data. Initial studies of subsets of the data allowed for comparisons of the configuration of submodels used within the precipitation framework. Once a final framework configuration was selected, uncertain interfacial tension parameters were optimized against the experimental dataset, and the model's performance at capturing the full range of available data was evaluated. While the framework was not able to capture long time-scale trends well, it performed satisfactorily at capturing short time-scale trends across a wide range of temperatures. Additional types of framework data outputs were shown as examples of datatypes that would be available if the framework was applied to a different system and had access to other types of data. Finally, the sensitivity of the framework to the uncertain interfacial tension and mixing time parameters was illustrated, demonstrating the necessity of greater emphasis upon model validation and uncertainty quantification for future work on similar frameworks.

CHAPTER 4

CREDIBLE SIMULATION DEVELOPMENT

4.1 Introduction

It has long been recognized that when scientific phenomena are computationally modeled, an inherent need to quantify the uncertainties involved with such computations exists [124]. From this opportunity the validation and uncertainty quantification (VUQ) field of study has been defined. Through statistical analysis of experimental data, high-fidelity simulations, as well as low-fidelity simulations, greater confidence in simulations' ability to approach truth or underlying physics is obtainable. Modern VUQ encompasses topics such as experimental design [79], model calibration [74, 135], uncertainty quantification of model and experimental data [147], and model predictivity [61, 107]. VUQ methodologies are essentially an elaboration of the scientific process.

The following chapter will be organized as follows. First, descriptions of two VUQ methodologies are outlined: the National Institute for Statistical Sciences (NISS) [11, 12] approach and the consistency analysis [34, 33]. These two methodologies provide the backbone of the VUQ related research contained within this dissertation. Next, a process and theory for credible simulation development, which is a means of framing the scientific method for modern simulation development, is provided. Following this is an application walk-through of the credible simulation development cycle utilizing a spring-mass-damper pedagogical example.

4.2 Validation and Uncertainty-Quantification

Choosing an appropriate VUQ approach is an application specific activity and there is likely no single correct approach. Two VUQ approaches will be considered throughout this dissertation: a probabilistic method and one based upon basic engineering tests. The probabilistic approach describes uncertainties in terms of probability distributions, while the engineering method considers bounds. Where the probabilistic method applies

Bayes theorem to determine posterior probabilities, the other method utilizes data analysis techniques to locate regions in which the model is deemed consistent. The probabilistic approach considered is that of the NISS group, which is based largely upon the Kennedy and O'Hagan [74] Bayesian methodology, and the engineering test based method is known as consistency analysis [34, 122, 147].

4.2.1 NISS Approach

The NISS group presented a framework aimed towards the validation of computational modeling of physical processes. This framework was referred to as the Simulator Assessment and Validation Engine (SAVE) within Liu (2008) [86]. The framework is comprised of a six step recipe by which model VUQ could be approached. Bayesian statistical methodologies were the foundation of this framework. Initially, key sources of error and uncertainty within parameters and data are determined and quantified in an input/uncertainty (I/U) map. It could be argued that model-form uncertainty, as will be described in Chapter 5, could also be included in the I/U map due to being highly coupled with parameter uncertainty. From the I/U map, parameters are selected as active parameters for the current VUQ analysis based on sensitivity ranking and how many parameters the analysis can afford to investigate. Active parameters are the parameters whose uncertainty will continue to be considered throughout the investigation, while all other parameters are fixed to their nominal values. Bayesian methods are then utilized to propagate the active parameter's uncertainties through the computational model, allowing the model outputs to contain full statistical measures. Frequently, computations are expensive with respect to capital and time, so surrogate models are constructed over sparse parameter and model output data allowing interpolation to predict behaviors in regions of interest. Continuing with the Bayesian influence, Gaussian processes (GP) were a suggested type of surrogate model [11, 12]. Next, analyzing the comparison of the model and experimental data leads to gained knowledge, also known as validation. The gained knowledge can then be applied as prior knowledge for a new cycle of VUQ. This VUQ cycle can easily be mapped into the more generic credible simulation development cycle that will be presented within Section 4.3.

A Bayesian methodology upon which the the NISS VUQ approach was partially based and which can be used for a subsection of the overall NISS process is the methodology presented by Kennedy and O'Hagan (2001) [74], which is now commonly known throughout the VUQ community as the Kennedy O'Hagan method. Within this methodology GPs are used as surrogate models to represent simulations. These GPs are then compared with

experimental data, and the difference between the GP inferred values and experimental data is also known as the model discrepancy. This discrepancy is then modeled with GPs as well. This approach allows for parameter calibration and uncertainty propagation through its Bayesian probabilistic framework, as well as a function discrepancy that can be utilized to study model-form issues. Interpolation once the model's discrepancy is accounted for generally performs well, but extrapolation is still problematic. The original paper describing this methodology [74] has currently been cited over 1,200 times as of April 2015. A more in-depth look into the Kennedy O'Hagan approach is provided within Chapter 5.

4.2.2 Consistency Measures

Another approach towards VUQ known as consistency measures was developed within the combustion community by Michael Frenklach and Andrew Packard at the University of California Berkeley (UCB). This UCB group developed their VUQ methodology as part of an effort to characterize the full methane-combustion reaction scheme known as GRI-Mech. Consistency measures were first proposed within Frenklach (2002) [41] and then further developed and defined by Feeley (2004) [34]. The use of consistency measures, implemented within a Matlab code published as Data Collaboration, was then demonstrated in model comparison capacities within Feeley (2006) [33]. Frenklach (2007) [40] presented the philosophy built within their consistency analysis tool, now contained within a framework called PrIMe, as a new scientific method of approaching model predictivity. An explanation and exploration of sensitivity analysis of the parameter and error uncertainties was then presented in Russi (2008) [121]. Optimizing the information contained in descriptions of consistent parameter spaces was explored by Russi (2010) [122]. You (2011) [147] reported upon the use of several different optimization schemes for defining the consistent parameter set. Two PhD dissertations stemmed from the development of this consistency analysis [32, 120]. The PrIMe tool was recently utilized by another research group for uncertainty quantification of a soot particle model [125]. The UCB group currently refers to this method as bounds to bounds consistency, or b2b.

The basic concept of consistency analysis boils down to comparing modeling outputs with experimental data. A simplified version of this concept can be described mathematically by

$$\begin{aligned} l_i \leq M_i(\mathbf{x}) - d_i \leq u_i \qquad & \text{for } i = 1, ..., N \\ \alpha_p \leq x_p \leq \beta_p \qquad & \text{for } p = 1, ..., n\,. \end{aligned} \tag{4.1}$$

Here the difference between model output M and experimental data d is compared to the error bounds of the experimental data l and u. This comparison can be conducted for

N data points, which can include multiple types of data. Comparing multiple data types allows for the model-form to be more strenuously tested. The other major component of the consistency analysis is the refinement of the parameter space $\mathbf{x}$ composed of n parameters, which initially are bounded by α and β. If consistency comparisons are simultaneously met for all data points using the same parameter subspace, the model is deemed consistent. This methodology's utility is not limited to assessing overall consistency. A great deal more information about the model and experimental data can be gained during the analysis through unary consistency, binary consistency, sensitivity to individual data points, as well as parameter sensitivity.

4.3 Credible Simulation Development

The scientific method has long been utilized as a means of progressing scientific thought. An ever increasing amount of science is now being undertaken through computational simulations. Here, I wish to provide a general description of a conception of credible simulation development, as currently conceived within the University of Utah's Carbon-Capture Multidisciplinary Simulation Center. Credible simulation development can be thought of as a means of framing the scientific method in terms of modern simulation science. The credible simulation development cycle currently utilized within the CCMSC is depicted within Figure 4.1.

Previous visions of this topical area were discussed within many standards documents such as those generated by the American Society of Mechanical Engineers (ASME) and Sandia National Laboratories [136, 105, 8]. Such standards were typically based fundamentally upon the Sargent Circle [126], which shares many similarities with the proposed credibility cycle. Where these two methodologies differ are the emphasis on the separation

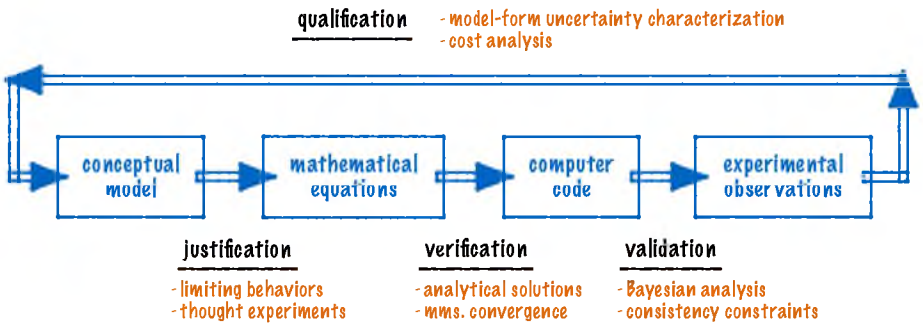

Figure 4.1: Cyclic process for developing credible simulations.

of the conceptual model and mathematical equations, additional justification process, and how the qualification process is handled. Although previous descriptions of the cycle have separated the creation of conceptual models and mathematical equations, this separation was approached as an elaboration on the Sargent Circle approach and not considered a fundamental piece of the process as it is within this methodology. First, a brief description of all blocks and connecting processes comprising the credible simulation development cycle will be provided. Following those characterizations will be an application walk-though.

4.3.1 Conceptual Model

The conceptual model is the current theoretical understanding on a phenomena of interest. It is generally built upon prior knowledge of the phenomena of interest or phenomena that appear to be similar in nature. When constructing a conceptual model the intended application is identified and areas of emphasis are ranked. Conceptual models are composed of the current, complete vision of all pieces of the phenomena, with no portions ignored due to preconceived irrelevance. Amassing the full description of the phenomena in this step allows for assumptions of the perceived importance of individual details to be weighed and compared in future steps, but all potential information is available for future reevaluations.

4.3.2 Mathematical Equations

The best representation of the conceptual model in terms of understood mathematics is what is being referred to with the mathematical equations step of the credible simulation process cycle. This step often builds heavily upon prior knowledge about the sorts of mathematical equations that have previously been used for similar models. Often a system of equations with initial and or boundary conditions are necessary to specify the conceptual model in a manner satisfactory for the conceptual model's intended use. The set of equations selected are designed to suite the application's characteristics of interest, not necessarily the full conceptual model. A continuum phenomena may not need an atomistic description and the same is true of the reverse situation. Creating a model for the appropriate model-scale is a practice known as scale-bridging.

4.3.3 Computer Code

Both the computational language and the algorithm used to numerically mimic the mathematical equations are considered parts of the computer code step. The specific computer language utilized will determine how solving the equations of interest can be structured. Algorithms translate mathematical equations into computational implementa-

tions that can be solved with a variety of numerical methods. To ease the implementation of numerical solutions or algorithmic needs, the form of the mathematical equations might be altered. An important part of this step is the selection and implementation of appropriate numerical methods. Different numerical methods will cause different amounts of numerical error to be injected into the solutions of the mathematics.

4.3.4 Experimental Observations

Experimental observations are just that, recorded experimental data with which the conceptual model can be compared via outputs of the computer code. Sometimes high-fidelity model data can be used as a substitute for experimental data, but this is not optimal. There is generally some error within the experimental data and exactly what the recorded data actually represent needs to be considered. Although conceptually experimental data are simply observations of reality, in practice, instrumentation is typically utilized to process raw data into reported data. How such instrumentation functions and potentially how the instrumentation was actually utilized often needs to be modeled, effectively creating a secondary credible simulation loop. Models describing how experimental data were collected can be referred to as instrument models. Even with many potential types of error within the experimental data, it is still generally held as being our best means of estimating the underlying 'reality' that our models are aimed to mimic.

4.3.5 Justification

The process of evaluating the mathematical representation of the conceptual model can be referred to as justification. When selecting the appropriate pieces of mathematics believed to represent the contrived concept, the selection process should be based upon limiting justifying tests. Such tests can include searching limiting-behaviors and considering thought-experiments. If a mathematical component passes all tests currently believed to be important for the application of interest, then the mathematical model has proven worthwhile for moving forward. If not, a different mathematical representation needs to be devised.

4.3.6 Verification

When translating mathematical equations into a computational language many error-checks and characterizations are necessary. ASME defines verification as 'the process of determining that a model implementation accurately represents the developer's conceptual description of the model and the solution to the model' [136]. Useful tools to aid in this

translation process include analytic solutions and method of manufactured solutions convergence studies. Examples of errors that such comparisons can illuminate include transcription errors, algorithmic issues, and discretization based errors. Each error type can be dealt with differently. For example, discretization error, or the error due to numerically solving mathematical equations across discrete space instead of a continuum, can be characterized with Richardson extrapolation during postprocessing [104]. It is during this step that the error created by the mathematical modeling technique and implementation thereof can be quantified and this knowledge carried forward. As long as errors are recognized and quantified for any computed outputs, they should not cause significant issues moving forward.

4.3.7 Validation

Quantifying the computational model's ability to reproduce experimental data is known as validation. Within validation processes lie opportunities for parameter calibration. By comparing model outputs produced while varying parameter values within their defined uncertainty against experimental data, the range or distribution of likely parameter values can be learned. Likewise, model outputs can be compared against other experimental data to ascertain regions of model validity. Measures of validity are often relative and many methods of approaching such definitions exist [98, 104]. An example of a validation metric is the area validation metric wherein the cumulative distribution function of model outputs is compared with an empirical cumulative distribution function for experimental data [98]. It should be stressed that data sets used for calibration and validation should not overlap in order to ensure confidence in model validity statements. An additional task that can be achieved throughout this step is uncertainty quantification, whereby pushing parameter uncertainties through the model into outputs allows uncertainty in the model predictions to be quantified. Two means of accomplishing these validation tasks are Bayesian analysis [74] and testing prescribed consistency constraints [34], as were previously mentioned in Section 4.2.1-4.2.2. Accomplishing calibration and uncertainty quantification during validation provides greater understanding of the model's true predictive capabilities than would otherwise be available.

4.3.8 Qualification

After a model's validity has been assessed, a credibility decision known as qualification should be considered. Qualification determines if the current state of the model creates predictions adequate for the application. This analysis can include characterization of the

discrepancy between the model predictions and experimental data due to the model-form, also known as model-form uncertainty. All considerations here are framed within a cost analysis mindset, where the cost-benefit of completing an addition cycle of model refinement must be evaluated. Depending upon the severity and character of the discrepancy between the model prediction and experimental data, the amount of time taken in reconsidering the conceptual model versus altering the mathematical formulation or algorithmic implementation utilized can be determined in order to most efficiently progress.

Idealistically, there is a clear distinction between model validity and adequacy. Model validity is meant to characterize a model's ability to represent the 'truth,' while model adequacy quantifies the model's capacity to satisfy criteria specified by the practitioner. Clearly, adequacy is application specific, but through successive iterations of the credible simulation development cycle, model adequacy should converge towards model validity. It would seem that as long as discrepancy between model prediction and experimental data exists due to model-form error, only adequacy and not validity can be claimed.

4.4 Application Example

For an application walk-through of the credible simulation development cycle, a spring-mass-damper system will be utilized (Figure 4.2). This system involves a particle of mass m attached to a spring and a damper. For this example the behavior of the particle once it has been pulled away from an equilibrium position, flicked further away, and then evolves through ever lessening oscillations for 20 time units is what is meant to be simulated. Our conceptual model of this phenomena stems from our understanding of Newton's second law, or that mass times acceleration $\ddot{x}$ of a body is equivalent to the sum of forces acting on the body,

$$\begin{aligned} m\ddot{x} &= \sum \text{forces} \\ &= f_b - f_d - f_s. \end{aligned} \tag{4.2}$$

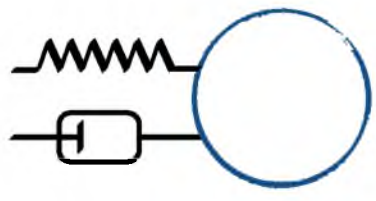

Figure 4.2: Illustration of pedagogical example spring-mass-damper system.

For a spring-mass-damper system the forces acting upon the particle will be a spring force f_s, damper force f_d, and general body force f_b that initiates the system evolution by moving the particle from its equilibrium position. Due to our prior knowledge of springs, it is known that the spring force will pull the particle towards its equilibrium position with force linearly proportional to the amount the spring is stretched, also known as Hooke's law of elasticity. Likewise the damper force will act against movement of the particle based upon the velocity of the particle. The body force for this example will only effect the particle initially by providing the initial conditions from which the particle's evolution is tracked.

This system can further be considered in terms of energy balance. It can be conceived that once the particle is released, the spring will pull the particle and the damper will resist the particle's movement. Without the damper the particle would infinitely oscillate as the spring's energy was converted from potential energy to kinetic energy and then back into potential. Instead of this infinitesimal oscillation, the damper absorbs some of the kinetic energy, effectively siphoning away a portion of the energy being passed between the particle and the spring. Then once the energy is reduced to an amount not substantial enough to move the particle, the system will settle to rest in the particle's equilibrium position. This description in terms of forces and energy comprise the conceptual model that will be carried forward for this cycle of simulation development.

Next, mathematical equations need to be created to describe the conceptual model. Linear relations will be utilized to describe the spring and damper forces,

$$f_s = -kx \qquad f_d = -c\dot{x}. \tag{4.3}$$

These two mathematical definitions can be tested with thought-experiments or limiting-behaviors based upon the conceptual model. Given our conceptual understanding of springs, the maximum amount of spring force should be available when the spring is at the extreme points of an oscillation and the minimal spring force should occur at the equilibrium position. The devised relation for the spring force is linearly proportional to the particle position causing greater amounts of force to be applied as the distance from equilibrium increases and no force occurs at the equilibrium position. Thus, this liner relationship fits our limiting behaviors. Conceptually, our understanding of the damper is less well refined than the spring, but we do understand that the damper force is related to the amount of kinetic energy available, which is based upon the particle's velocity. Having the damper relation be linearly proportional to the particle's velocity causes maximum dampening for each oscillation to occur at the equilibrium position where all spring energy have been converted to kinetic energy. Minimal dampening then occurs at the maximum positions of oscillation

where the velocity is zero and no kinetic energy remains. This linear relationship again aligns well with the understood limiting-behaviors. Both the spring and damper forces have a negative sign assigned to their relations in order to match the orientation of positive position and velocity relative to the equilibrium point that is assigned coordinate position zero.

For this example the body force will be encapsulated as initial conditions for the mathematical equations. To describe moving the particle a distance away from equilibrium and then flicking it further from equilibrium as a release mechanism for the particle, an initial position of six and an initial velocity of three can be utilized. A particle mass $m = 3$ and a spring coefficient $k = 5$ will also be assigned as known values, but the damper coefficient will be set as an unknown parameter with Jeffreys prior distribution spanning from zero to $2\sqrt{km}$ to reflect uncertainty in our understanding of the particular damper being utilized in the system, as well as our previously stated less refined understanding of dampers. This uncertainty was recognized as the potential range of values because the system was characterized as falling within underdamped conditions within the conceptual model, or an oscillating system. Mathematically it can be shown that for this system to be underdamped it must meet the condition that $c < 2\sqrt{km}$ and a damper coefficient of zero would remove the damper from the system [78]. A summation of the mathematical description of the conceptual model can be expressed as

$$\begin{aligned} &x(0) = 6.0 \qquad \dot{x}(0) = 3.0 \\ &m\ddot{x} + c\dot{x} + kx = 0. \end{aligned} \tag{4.4}$$

To convert this mathematical equation into a computational algorithm, the second order ODE can be converted into two first order ODEs

$$\begin{aligned} \dot{v} &= -\frac{k}{m}x - \frac{c}{m}v \\ \dot{x} &= v, \end{aligned} \tag{4.5}$$

also known as the state space representation. Eq. (4.5) can then be temporally discretized with Euler's method to yield

$$\begin{aligned} v_{N+1} &= v_N - \frac{k}{m}x_N\Delta t - \frac{c}{m}v_N\Delta t \\ x_{N+1} &= x_N + v_N\Delta t, \end{aligned} \tag{4.6}$$

where N and $N+1$ differentiate the current and next temporal positions and Δt symbolize the time-step. These equations can then be solved with a numerical time integration method

initiated from time zero with the specified initial conditions. To solve these equations over the desired time domain of 20, a predictor-corrector variety of an adaptive Heun-Euler time integrator can be used with two corrector cycles. Eq. (4.6) and the numerical time integrator have been coded as Python scripts. To verify the numerical implementation, an analytic solution to this system of equations is available [78]. The left plot of Figure 4.3 compares the analytic and numerical solutions of the spring-mass-damper system and the right plot shows the discrepancy between those two solutions. The discrepancy between the numerical and analytic solution indicates that the numerical approximations utilized cause up to 0.0002 absolute error but appear stable.

Now to compare against experimental data, synthetic experimental data will be generated from an example 'truth' model presented by Oliver et al. (2015) [107]. The underlying 'truth' for this case was that the damper coefficient was temperature dependent and the temperature of the damper was effected by the velocity of the particle. This relationship tries to mimic the damper fluid heating due to absorbing the particle's kinetic energy and then cooling as the fluid equilibrates with the surroundings via heat transfer. The damper coefficient and temperature changes are described by

$$c(T) = \exp\left(\tfrac{T_0}{T} - 1\right) \tag{4.7}$$

$$\dot{T} = c(T)\dot{x}^2 - \tfrac{1}{\tau}(T - T_0), \tag{4.8}$$

where $T_0 = 20$ is a reference/initial temperature and $\tau = 1$ is a heat transfer scaling factor. Synthetic data were created by taking samples along the 'truth' model's position trace and

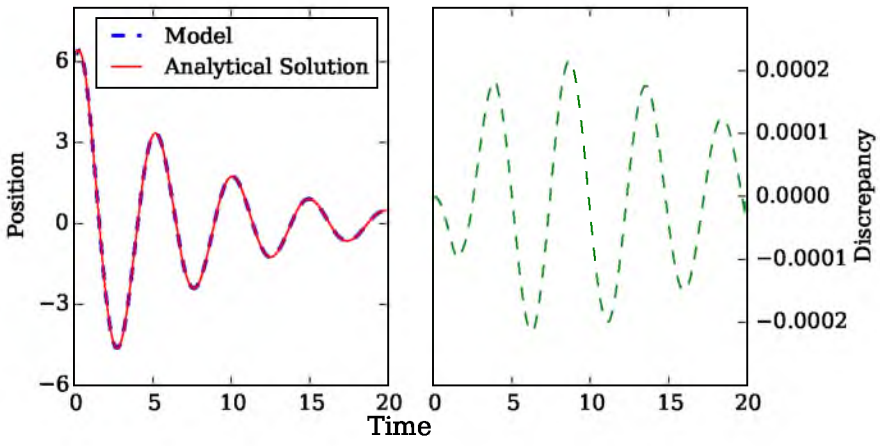

Figure 4.3: Verification of numerical implementation of Eq. (4.6) solved with an adaptive Heun-Euler time integrator by comparing with an analytical solution. Left plot shows traces of analytical (red line) and numerical traces (dashed blue line) over time and right plot shows the discrepancy between the solutions over time.

adding random Gaussian noise with a standard deviation of 0.1, $\mathcal{N}(0, 0.1)$. Three data replicates were created for each position and these data points are shown in the left plot of Figure 4.4.

The right plot of Figure 4.4 demonstrates the effect of the uncertainty in damper coefficient value. Traces in that plot were generated by uniformly taking five samples from the variable's potential values between 0 and 7.45. Clearly the prior uncertainty in the damper coefficient causes significant uncertainty in the model outputs when propagated through the model.

The next step is to perform calibration on the uncertain damper coefficient. Two calibration methods will be illustrated: an engineering consistency approach and a Bayesian parameter estimation. These methods are not exactly the two VUQ approaches discussed within Sections 4.2.1-4.2.2, but are similar and provide a feel for the differences between those methods. For both methods the prior range of the calibration parameter, 0 to $2\sqrt{km}$ will be explored.

For the consistency approach the prior knowledge about the calibration parameter will be utilized as prior bounds. Then sampling parameter values across that prior range, position traces are calculated. These position traces are compared with the experimental data and consistency can be assessed. First L_1 norms are calculated for each temporal position that experimental data are available for,

$$L_{1,x} = \sum_{i=n_x} \frac{|y_{M,x} - y_{E,x,i}|}{n_x}. \tag{4.9}$$

Once L1 norms exist for each experimental temporal position an L_∞ across the L_1 norms can be calculated, Eq. (4.10), and this value is then compared to the desired consistency constraint,

$$L_{\infty,t} = \max(L_{1,x,t}). \tag{4.10}$$

Figure 4.5 illustrates how this method was applied to calibrating the current system with the left plot demonstrating the performance across the entire prior parameter range and the right providing a zoomed-in view of the better performing region.

Bayesian parameter estimation is accomplished using Bayes law Eq. (4.11) where the prior distributions $P(c, \sigma)$ for unknown parameters must be specified, as well as a likelihood function $P(y|c, \sigma)$ for the data. Here, we assume that the noise in the data is Gaussian distributed, so that the data were distributed as $\mathcal{N}(y_E, \sigma)$, and through a variable transformation the likelihood also takes a Gaussian form, Eq. (4.12). For the calibration

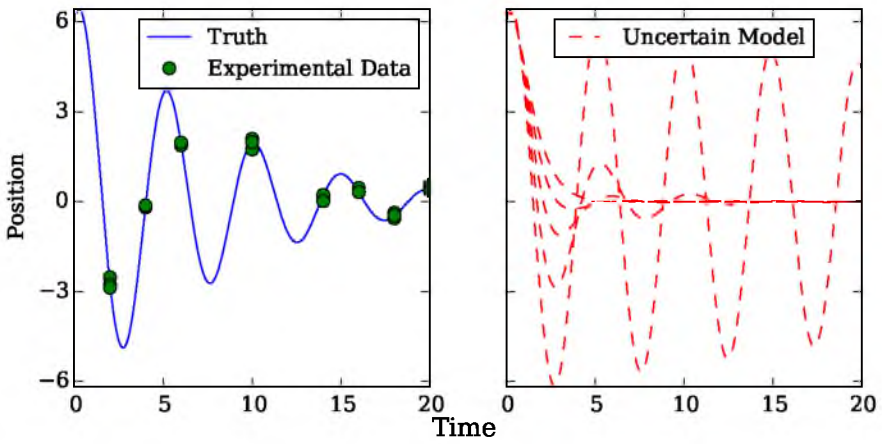

Figure 4.4: Comparing underlying true physics, experimentally observations, and uncertainty in model thereof. Left plot displays 'true' model trace and experimental data points synthetically created by adding random samples of Gaussian noise with a standard deviation of 0.1. Right plot shows traces from model Eq. (4.6) containing uncertain damper coefficient where the different traces correspond to samples of the uncertain parameter's initial uncertainty range between 0 and $2\sqrt{km}$.

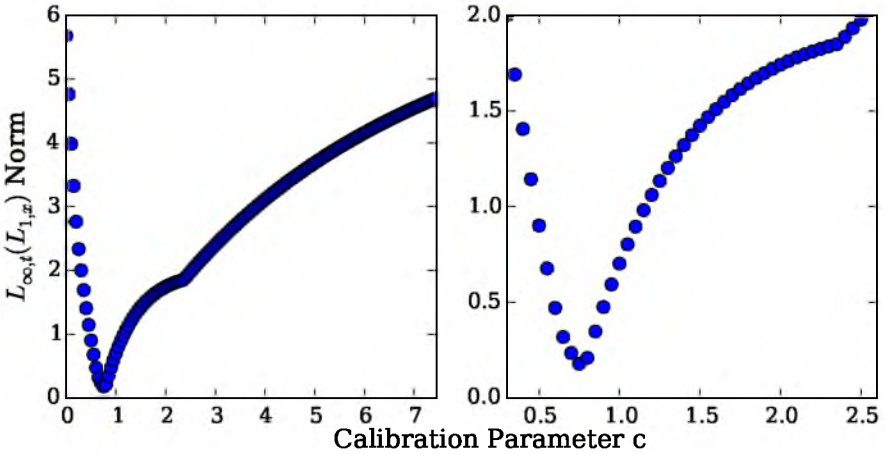

Figure 4.5: Parameter calibration of uncertain damper coefficient c based upon consistency constraints. Left plot displays the error between the experimental data and model-form Eq. (4.6) across the prior range of calibration parameter values. Error is calculated as the L_1 norm across experimental data replicates and then L_∞ norm of the resulting values across temporal positions. Right plot is a zoomed in view of left plot, where the region of likely parameter values lies depending upon desired consistency constraints.

parameter a bounded Jeffreys prior can be utilized as an uninformed prior for the parameter range. With the Bayesian approach the standard deviation of the data's noise must also be considered and a student's t-distribution will be used as the prior, but we assume this distribution will converge to Gaussian if enough data points are available. By assuming the prior distributions are independent, they can joined as a joint prior $P(c)P(\sigma) = P(c,\sigma)$. This can be expressed as

$$P(c,\sigma|y) = \frac{P(y|c,\sigma)P(c,\sigma)}{\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} P(y|c,\sigma)\mathrm{d}c\,\mathrm{d}\sigma} \tag{4.11}$$

$$= \frac{1}{(2\pi\sigma^2)^{n_{x+t}}} \exp\left(\frac{-\frac{1}{2}SSq}{\sigma^2}\right) \cdot \frac{P(c,\sigma)}{\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} P(y|c,\sigma)\mathrm{d}c\,\mathrm{d}\sigma}, \tag{4.12}$$

where

$$SSq = \sum_{i=n}(y_{M,i} - y_{E,i})^2.$$

The denominator of Bayes law does not need to be explicitly solved here because the fact that the posterior distribution integrates to one can be exploited. Figure 4.6 shows the posterior (left) and posterior with the noise's standard deviation marginalized over (right).

Clearly the consistency and Bayesian methods produce different results, but the ultimate conclusions of the two methods are similar. For this example the bound for consistency will be set as less than 0.30 based upon the amount of spread that could be observed in the experimental data. Figure 4.7 demonstrates that the consistency method locates a calibrated region of parameter space between approximately 0.65 and 0.82, while the

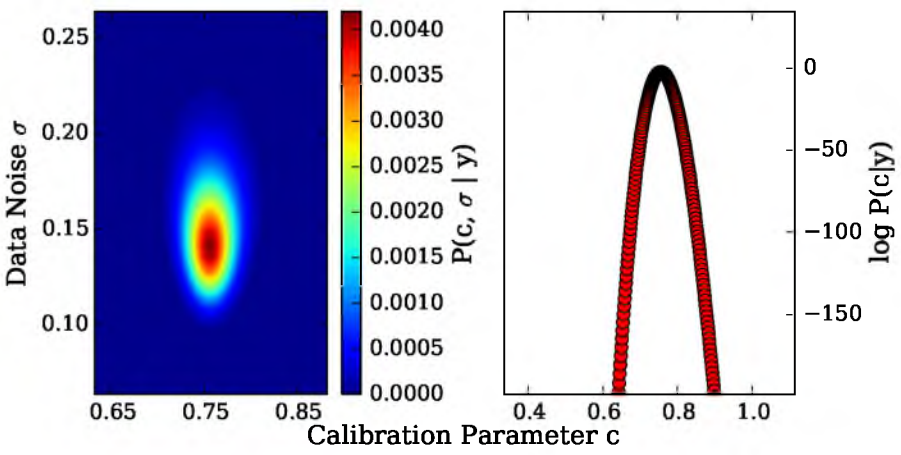

Figure 4.6: Parameter estimation based upon Bayes law for uncertain damper coefficient c and noise standard deviation in experimental data σ. Left plot displays posterior distribution $P(c,\sigma|y)$ across ranges of c and σ. Right plot shows marginal posterior distribution $\log P(c|y)$ across a range of c values.

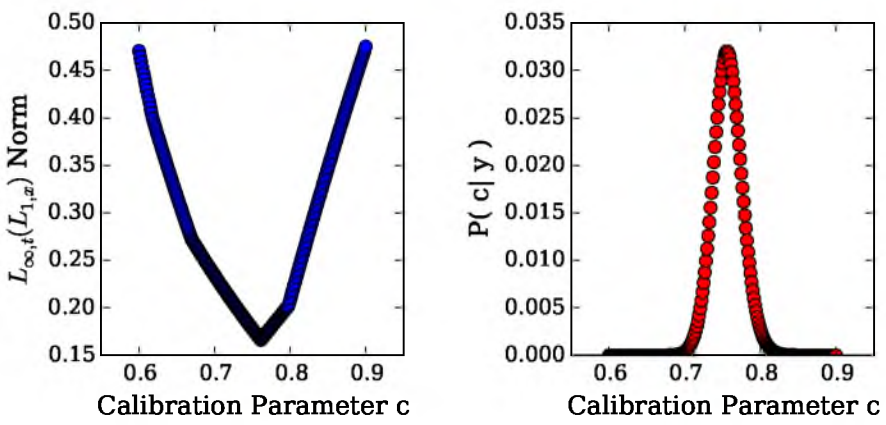

(a) (b)

Figure 4.7: Comparison of calibration result. Plot (a) shows consistency constraints and and plot (b) shows Bayesian parameter estimation.

Bayesian estimate also locates a similar region between 0.71 and 0.81. The most consistent point occurs when the calibration parameter was 0.76 and the most probable point with the Bayesian method is at 0.75. The most probable standard deviation of the experimental data noise was 0.14. The major difference between these methods is that the Bayesian method provides a distribution estimate, while the consistency method provides bounds. Ultimately, this difference will result in different types of characterizations of uncertainty in predictions made by the model.

Now that calibrated parameter regions have been identified, validation of the current model-form can occur. First, the uncertainty of the calibration parameters can be propagated through the model, so that the uncertainty in the model's predictions can be quantified. Where calibration can be referred to as the inverse problem, uncertainty propagation is then the forward problem. Figure 4.8 shows the uncertainty present in model outputs due to the forward uncertainty propagation, as well as illustrates that the parameter calibrations performed well at capturing the experimental data. The consistent region, which is shown as the grey region, represents the full parameter region found to be consistent, while two standard deviations from the probability based approach is shown as the red region where 50 samples were taken form the marginal distribution shown in Figure 4.7b. Both methods produced approximately the same best estimate of the trace, but their uncertainties have different meanings. Where the Bayesian method produces a distribution of uncertainty across time, the consistent method produces an uncertain region across time.

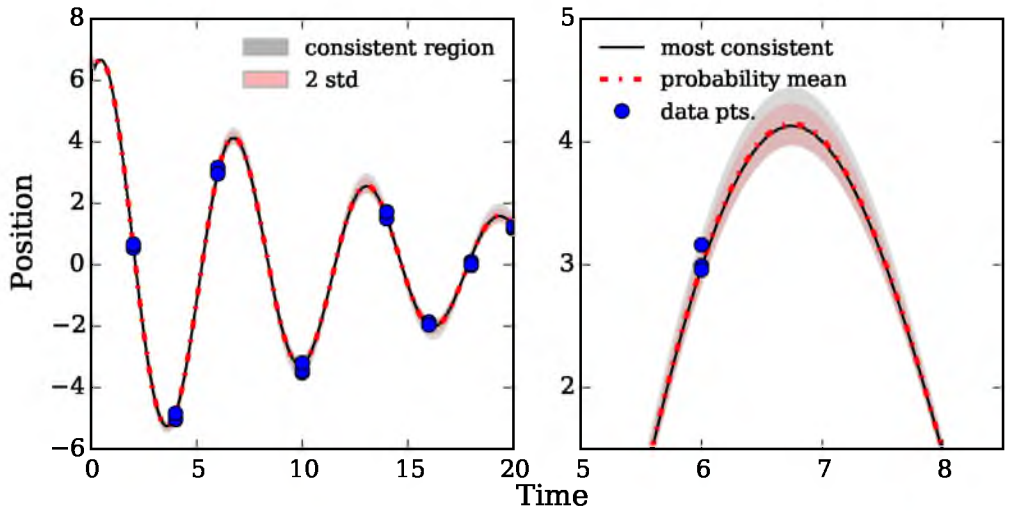

Figure 4.8: Forward propagation of parameter uncertainty through model into model outputs. Here blue dots represent validation data points, the black line is the model run with the most consistent parameter value, the grey region represents the full uncertainty of potential model values based upon the full consistent parameter region, the red dash-dotted line is the model run with the most probable parameter value, and the red region represents two standard deviations of the uncertainty in the model prediction based upon the calibrated probability distribution. The left plot shows the full temporal range and the right plot shows a zoomed in view of a section of the left plot.

To validate the model an experimental data set, distinct from the data used for calibration, must be available. Thus, for validation another synthetic data set can be created with Eq. (4.6) using the temperature dependent damper coefficient Eq. (4.7) and evolving the damper's temperature with Eq. (4.8). What distinguishes this data set from the calibration data is that the particle mass is set to five instead of three. This change in particle mass effectively represents changing the experimental system and should test if the calibration had a physical basis, or due to an incorrect model-form only fit to the data. The validation dataset and model performance is demonstrated within Figure 4.8.

To visualize the validity of the current model-form, discrepancies between the experimental data and model predictions are plotted within Figure 4.9. Plots (a) and (b) demonstrate the effect upon the model validity of using the consistency versus the Bayesian method of calibration. Within these plots the optimal calibration value for each method is plotted, as well as the full calibration range. The dots indicate mean values and the error-bars show ranges. For the Bayesian parameter estimation, the distribution shown in Figure 4.7b had 50 random samples taken to reflect the consistent region. Plot (c) demonstrates the discrepancy between the data and model for the calibration data set, illustrating the models optimal performance for the calibrated parameter values. Evidently, optimal values found

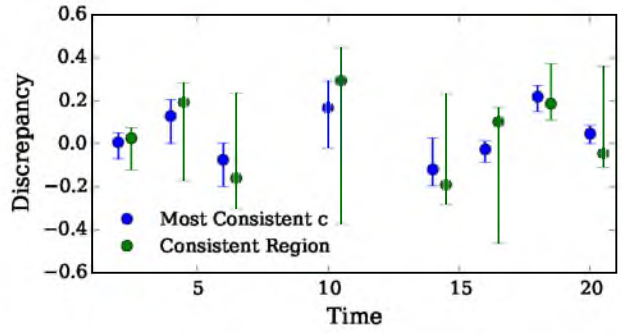

(a)

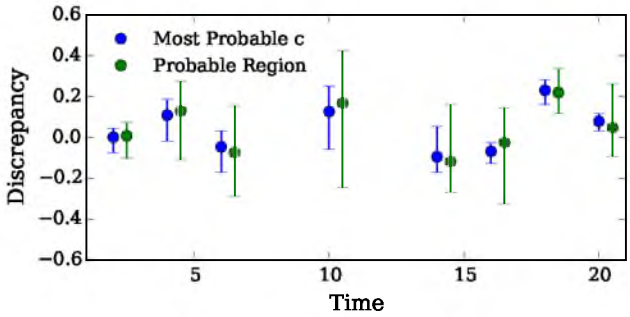

(b)

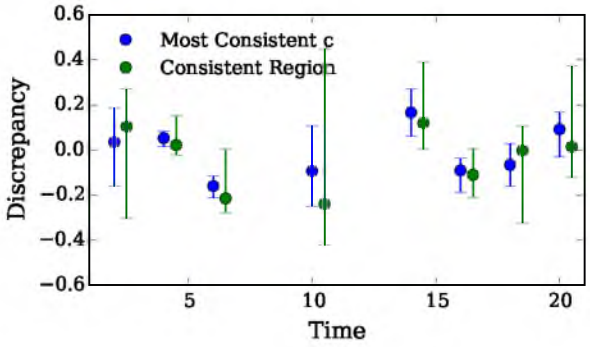

(c)

Figure 4.9: Comparison of discrepancies between model trace outputs and validation experimental data traces where the data were spread over the temporal domain. For all three plots the dots indicate mean discrepancy values across the data replicates. Plot (a) shows the discrepancy when the consistency calibration results are used, where the green error-bar is the most consistent point and the blue error-bars represent the full range of consistent calibration values. Likewise, plot (b) displays discrepancy when the Bayesian calibration is utilized, where the green errorbars reflect the most probable calibration value and the green error-bars correspond to 20 random samples from the calibration's marginal posterior distribution, Figure 4.7b. Plot (c) is equivalent to plot (a) except that the traces were compared to the calibration data set.

with both calibration techniques perform well and have mean discrepancies within ± 0.2. By comparing the discrepancy based upon the validation data set against discrepancy based upon the calibration data set, it appears that the quantities of discrepancy are similar for the model with both data sets. It seems that the calibration focused upon reducing discrepancy for the second and third temporal positions, but this effect did not translate to the validation system conditions. The conclusion of this validation step is that within or near the validation and calibration conditions the model should perform predictions with errors similar to those found during validation. The effect of the calibration method utilized appears to have minimal repercussions on this model's performance for validation.

The final step of the credible simulation development cycle is to access the qualification of the model-form. The first objective of qualification of the spring-mass-damper model involves determining if the model-form error, as shown in Figure 4.9, is significant enough to merit further investigation and ultimately an additional cycle of the model development process. For the sake of this example, the discrepancy will be deemed unacceptable. Next, the discrepancies can be further examined with the aim of locating issues that could be traced back to the model-form. One notable feature shown in Figure 4.9 is the increased amount of discrepancy for the second and third (in temporal order) data-points for the validation data set as compared with the calibration equivalent. Hypothetically, through continued examination it could be considered that this region of increased discrepancy occurs where the particle is transferring the most energy to the damper. This could lead to conceptual contemplation about what that energy does to the damper and the damper's temperature could be experimentally monitored. Figure 4.10a shows what experimental observation of the damper's temperature could uncover, based upon Eq. (4.8). A conceptual model for the temperature's effect upon the damper coefficient could then be constructed during the next development cycle.

Based upon Eq. (4.7), Figure 4.10b shows the 'true' damper coefficient's value for the calibration data set. It can be noted that the calibrated parameter ranges corresponded with the true damper coefficients for the initial three data point temporal positions. The optimal calibrated parameter values lie near the middle of the true range. This is not surprising, but indicates that the calibration methods performed as well as the model-form error allowed.

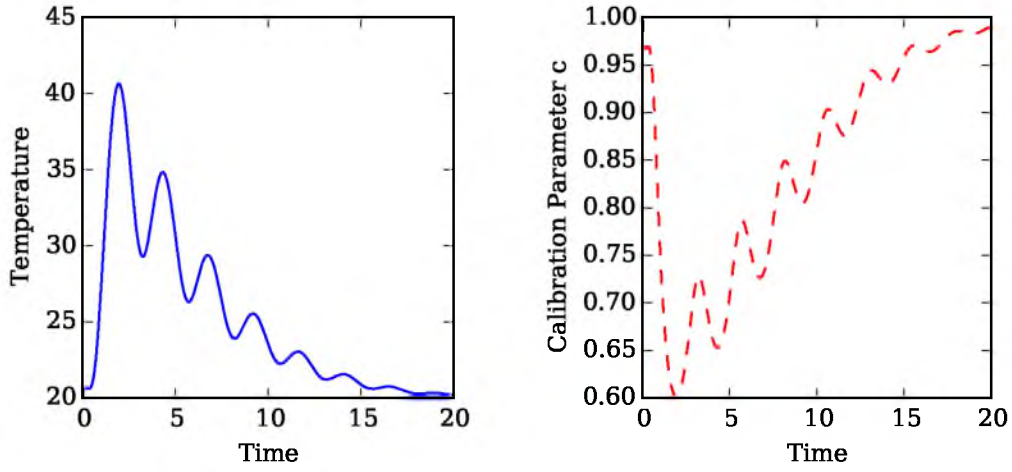

(a) (b)

Figure 4.10: Demonstrating 'true' effect of temperature on damper. Plot (a) shows temperature profile according to 'truth' model Eq. (4.8) underlying calibration experimental data set. Plot (b) then translates the effect of the damper's temperature on the damper coefficient according to Eq. (4.7).

4.5 Conclusions

Validation and uncertainty quantification is a broad area of study devoted towards improving performance and understanding of models. Many potential VUQ approaches exist but two, Bayesian and consistency analysis, were introduced and will be explored throughout the subsequent chapters. A process for developing credible simulations has been presented: consisting of conceptual model development, translation to mathematical equations, checking that transition with justification tests, coding the mathematics into algorithms, double checking the code with verification techniques, amassing experimental data, validating the modeling framework against the experimental data, and ultimately qualifying if the current model development is adequate for the desired application. By walking though a mass-spring-damper pedagogical example, a general feel for the process has been provided. Applying such thorough methodologies should produce a better product, as well as likely help debugging unforeseen consequences of modeling decisions.

CHAPTER 5

BAYESIAN MODEL-FORM UNCERTAINTY EXPLORATION

5.1 Introduction

Often in modern scientific research imperfect models are continually evolved through comparison with experimental observations. Although these models are known a priori to be imperfect throughout such comparisons, this is frequently not explicitly accounted for and any error associated with this fact is distributed into parameter uncertainty and/or experimental error. Explicit characterization of the extent of such model imperfection may allow for more appropriate use of the model, as well as extend the ability to use the model in predictive capacities. This new functional error quantity is known as model-form uncertainty, discrepancy, or model-bias. Calibration of model parameters within imperfect models is a key issue that must be considered. Bayesian statistical methods are one family of methods commonly utilized to approach model-form uncertainty problems.

The following validation exercise utilizes a pedagogical example provided in Bayarri et al. (2005) [11]. Within the example a method for approaching model-form uncertainty developed by Kennedy and O'Hagan (2001) [74], which was also utilized by Bayarri et al. [11], will be explored. Overviews of the Bayesian statistical philosophy, Gaussian processes, the Markov chain Monte Carlo technique utilized, as well as the Kennedy O'Hagan methodology and the implementation algorithm used within the current exercise will be presented first. Following that background information will be a thorough exploration of implementing this approach through a chemical-reaction based pedagogical example. This exploration will include studying calibration estimation techniques, sensitivity to priors, potential short-comings due to assumptions, possible information gain from the discrepancy functionality, system sensitivity analysis, the identification problem, multi-input Gaussian processes, Bayesian model-comparison, and constrained Gaussian processes. Lastly, use of the calibrated parameters in forward propagation of uncertainty is considered.

5.1.1 Bayesian Theory Overview

Bayesian statistical methods are derived from logic and probability theory. Sivia (1996) [128] and Gregory (2005) [53] provide basic overviews of Bayesian methods, while Jaynes (2003) [66] delves deeply into the underlying Bayesian doctrine. Gelman et al. (2013) [44] presents more advanced Bayesian techniques such as Markov chain Monte Carlo (MCMC), regression models, and nonparametric models.

The backbone and namesake of Bayesian methods is Bayes law

$$P(H|DX) = \frac{P(D|H,X)P(H|X)}{P(D|X)}, \tag{5.1}$$

where P denotes a probability function, H a hypothesis, D gained data, and X prior information [66]. Verbally, Bayes law can be stated as the posterior $P(H|DX)$, or new knowledge, is gained by multiplying the prior $P(H|X)$, or old knowledge, by the likelihood $P(D|HX)$, probability based upon observations. This is all normalized by $P(D|X)$, sometimes known as the evidence or marginalized likelihood, which is necessary to force the integral of the probability function to equal one.

Application of Bayes law to scientific endeavors is intuitive considering that the scientific method is based on building upon the current state of knowledge to gain new knowledge. Bayes law can be derived using the product rule of probability theory and a joint probability distribution with simple rearrangement,

$$P(A,B) = P(B|A)\ P(A) = P(A|B)\ P(B). \tag{5.2}$$

The other key probabilistic theory definition relied upon heavily throughout Bayesian methods is the sum rule, which is shown here through both the continuous and discrete forms of the aforementioned evidence

$$P(D|X) = \int P(D|H,X)\ P(H|X)\mathrm{d}H \tag{5.3}$$

$$= \sum_i P(D|H_i,X)\ P(H_i|X). \tag{5.4}$$

Common uses of Bayesian theory include hypothesis testing, model selection, uncertainty quantification, and parameter estimation/model calibration. Using prior distributions for uncertain model inputs and prescribed likelihood functions, Bayesian methods allow for joint posteriors to be calculated, which contain the maximum amount of potential information gain for the system [61]. Such joint posterior distributions are often explored and marginalized using MCMC methods and thus this method will also be included within this analysis of approaching model-form uncertainty [74, 61, 144]. Surrogate modeling is often necessary

when approaching many of the aforementioned statistical tasks. Following the examples of Kennedy and O'Hagan [74] and Bayarri et al. [11], the Bayesian based surrogate models, Gaussian processes (GP), will be utilized for all surrogate modeling within the current approach. Once model parameters have been calibrated and uncertainty in inputs has been pushed through the model into predictions, the model can be validated against experimental data with methods such as Bayesian hypothesis testing [84, 98].

A major focus of this exercise and most applications where characterization of model-form uncertainty is sought is parameter calibration. This is also known as parameter estimation. While least-squares and other techniques for the minimization of weighed residuals can be utilized for estimating parameter values, Bayesian calibration techniques will be explored throughout this example instead.

5.1.2 Gaussian Process Overview

Gaussian processes are stochastic processes, or a sampling of a density function over a defined function space, where the joint distribution over any finite set of test points is multivariate Gaussian [17]. Provided a set of functional observations f evaluated over an i dimensional parameter space, a GP can be utilized to regress the underlying unknown function, which is now treated as a random function. In order to characterize a GP, a mean μ and a covariance matrix K are needed,

$$f|X \sim GP(\mu, K(X, X)). \tag{5.5}$$

Means can take functional forms and are often described as a set of weighed basis functions such as $h(X)^T\beta$ where $h(X)$ is a set of fixed basis functions and β the respective weighting of each, but for instances with no prior knowledge they are often set to a constant such as zero. It is important to note that GPs will devolve back to the prior mean when far from data-points. A multivariate normal distribution with mean μ and covariance K are the assumed prior distribution for f, which is then updated to a conditional posterior through incorporation of the function outputs, as shown in Eq. (5.5). GPs are considered nonparametric because they assume the distribution has infinite dimensions, i.e., theoretically the model expands infinitely as additional data X_i are added.

To regress or interpolate to desired points f_* the joint distribution takes the form

$$\begin{bmatrix} f \\ f_* \end{bmatrix} \Bigg| X \sim \mathcal{N}\left(0, \begin{bmatrix} K(X,X) & K(X,X_*) \\ K(X_*,X) & K(X_*,X_*) \end{bmatrix}\right). \tag{5.6}$$

Here $K(X, X)$ is the covariance matrix of parameter values of f, $K(X, X_*)$ and $K(X_*, X)$ are the cross covariance matrix between parameter values for f and f_* and its transpose,

and $K(X_*, X_*)$ is the covariance matrix for parameter values of f_*. The conditional mean m and covariance cov of the posterior predictive distribution at points X_* can be solved as [17]

$$m(X_*) = E[f_*|X_*, X, f] = K(X_*, X)\big(K(X, X)\big)^{-1} f \tag{5.7}$$

$$cov(X_*) = K(X_*, X_*) - K(X_*, X)(K(X, X))^{-1} K(X, X_*). \tag{5.8}$$

These conditioned properties are estimates of the desired point and the uncertainty associated with that quantity.

Many covariance kernels exist and appropriate choices are application specific. A general overview of available kernels can be found in Rasmussen and Williams (2006) [115]. The squared-exponential Eq. (5.9) and noise Eq. (5.10) covariance kernels are utilized within GPs for the current application,

$$k_{SE} = \omega \exp\left(-\frac{r_c^2}{2l^2} \right) = \omega \exp\left(-\sum_i \frac{r_{c,i}^2}{2l_i^2} \right) \tag{5.9}$$

$$k_{noise} = \lambda^2 I. \tag{5.10}$$

Here r_c is the Cartesian distance between input parameters, I is an identity matrix, i refers to input dimensions, and the other parameters are hyperparameters that will be discussed shortly. The squared-exponential kernel is commonly utilized in modeling physical phenomena due to being stationary, smooth, isotropic, and infinitely differentiable. There are options for how multiple input variables can be handled with the squared-exponential kernel, but the form shown above will be used for this exercise. Noise in the observational data can be accounted for with the noise kernel. If no noise is added, the regressed distribution's mean will go through all data-points. When present, noise can be added to the data covariance terms $K(X, X)$ in the above equations resulting in the following posterior predictive distribution [115],

$$\begin{aligned} f_*|X_*, X, y \sim \mathcal{N}\big(&K(X_*, X)\big(K(X, X) + \lambda^2 I\big)^{-1} y, \\ &K(X_*, X_*) - K(X_*, X)\big(K(X, X) + \lambda^2 I\big)^{-1} K(X, X_*)\big). \end{aligned} \tag{5.11}$$

Here y is used to signify the functional output values once noise, which is assumed to be normal and independent, has been included in the description of output values of f [117]

$$P(y|f) = \mathcal{N}(y|f, \lambda^2 I). \tag{5.12}$$

Small amounts of noise are sometimes added to a GP to help keep the covariance matrix well conditioned and to dampen out noise due to numerics. Addition of noise also differentiates regression from interpolation.

Covariance kernels have additional unknown parameters, known as hyperparameters, such as the ω, l and λ found in the aforementioned kernels. These hyperparameters will not necessarily have physical meaning, but do increase the dimensionality of the posterior distribution. The squared-exponential kernel's length-scale hyperparameter l does have a physical interpretation, larger values indicate that statistical correlation between data-points occurs over a wider range. Thus length-scales are relative sensitivity measures for functions with multiple inputs. Either the full distribution of each hyperparameter is carried through the GP to the final posterior distribution, which is known as the full Bayes method, or an approximation can be made. To explore the negative log posterior

$$-\log\ P\big(hp|X,y,K(X,X)\big) \propto \frac{1}{2}y^T\big(K(X,X)\ +\lambda^2 I\big)^{-1}y\ + \\ \frac{1}{2}\log\big|K(X,X)\ +\ \lambda^2 I\big|\ +\frac{n}{2}\log 2\pi - \log\ P(hp), \tag{5.13}$$

a Newton method based upon gradients

$$-\frac{\partial}{\partial hp_i}\Big[\log\ P\big(hp|X,y,K(X,X)\big)\Big] \propto \frac{1}{2}tr\bigg(\big(K(X,X)^{-1}-\vec{\alpha}\vec{\alpha}^T\big)\frac{\partial K(X,X)}{\partial hp_i}\bigg) \tag{5.14}$$

and the Hessian

$$-\frac{\partial^2}{\partial hp_i \partial hp_j}\Big[\log\ P\big(hp|X,y,K(X,X)\big)\Big] \propto \\ \frac{1}{2}tr\bigg(\big(K(X,X)^{-1}-\vec{\alpha}\vec{\alpha}^T\big)\frac{\partial^2 K(X,X)}{\partial hp_i \partial hp_j}\ - \\ \big(K(X,X)^{-1}-2\vec{\alpha}\vec{\alpha}^T\big)\frac{\partial K(X,X)}{\partial hp_j}K(X,X)^{-1}\frac{\partial K(X,X)}{\partial hp_i}\bigg), \tag{5.15}$$

where

$$\vec{\alpha} = K(X,X)^{-1}y,$$

allows the maximum a posteriori (MAP) estimate of the kernel hyperparameter (hp) values to be calculated, reducing the number of unknown parameters within the surrogate modeling.

The MAP method differs from the maximum likelihood estimate (MLE), another common method of estimating hyperparameters, in that the likelihood $P(y|X,K(X,X),hp)$ is multiplied by the hyperparameters' prior distributions $P(hp)$ to produce a quantity proportional to the posterior instead of only considering the likelihood. Prior distributions for the hyperparameters should reflect prior knowledge about the hyperparameters to the extent to which it is known. For instance, it is often the case that the hyperparameters will only take positive values and the general estimates of variance in experimental data noise

can also be easily produced. The MLE estimate is based upon the best fit to the data, which can be thought of as equivalent to the MAP method with improper uniform priors set equal to unity $P(hp) = 1$. Prior distributions used throughout this exercise include the log-normal Eq. (5.16) and Jeffreys Eq. (5.17)

$$P(x|\mu,\sigma) = \frac{1}{x\sigma\sqrt{2\pi}} \exp\left(\frac{-\big(\ln(x)-\mu\big)^2}{2\sigma^2}\right), \quad x > 0 \tag{5.16}$$

$$P(x) = \frac{1}{x}, \tag{5.17}$$

where μ and σ are the distribution's mean and standard deviation. Hyperparameter priors are assumed to be independent of one another and thus can be multiplied to create joint priors

$$P(\lambda)\ P(\omega)\ P(l) = P(\lambda,\omega,l). \tag{5.18}$$

The use of MAP or MLE hyperparameter estimates cause the GP to make optimistic variance estimates due to the loss of the hyperparameters' uncertainty or distribution. These methods are still commonly utilized due to their simplicity and computational speed.

Calibration can also be estimated through Eq. (5.13) with the inclusion of a prior for the calibration parameters. A calibration distribution can then be computed by marginalizing over the hyperparameters

$$\int P\big(u, hp|X, y, K(X,X)\big)\ \mathrm{d}hp = P\big(u|X, y, K(X,X)\big). \tag{5.19}$$

Of course MLE equivalents are possible. Methods of estimating the calibration distribution will be further explored within Section 5.3.3.

5.1.3 Gaussian Process Example

A short example of the use of GPs and some of the features they offers as a surrogate model is shown within Figure 5.1. This GP was conditioned upon seven data-points and used squared-exponential and noise covariance kernels. The intensity of the noise kernel λ was set to 0.1 causing the GP to regress among the data-points instead performing interpolation. The regressed point corresponding to independent variable $x = 0.2$ is shown as well as its inferred derivative. The estimated amount of error in the regression is shown as the blue shaded region, equivalent to one inferred standard deviation. Slopes of the data points upon which the GP was conditioned can also be estimated, as is shown for two data-points at $x = 0.1$ and $x = 0.6$.

A visualization of the covariance matrix contained within the example GP is shown in Figure 5.2. Here it can be noted that the diagonal components of the the matrix are

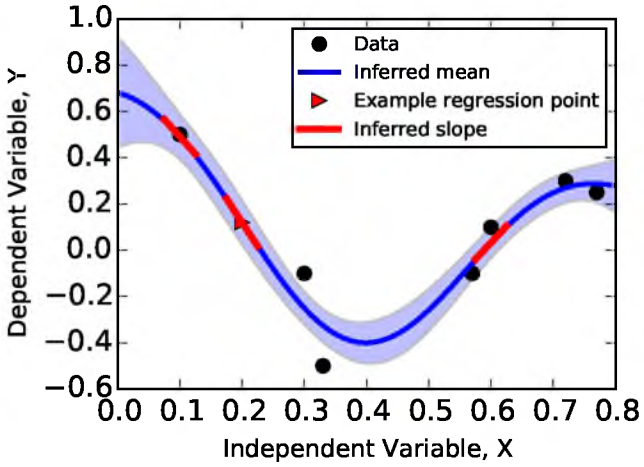

Figure 5.1: Example use of GP as surrogate model in a regression capacity. Initial data-points used as basis for GP are shown as black dots, an inferred data-point is shown as the red triangle, red lines are inferred slopes, the mean of the GP conditioned on the initial data is shown as the blue line, and the blue region represents one standard deviation of the conditioned GP.

the most correlated positions, corresponding to a data-point's covariance with itself. The diagonal values are not equal to unity due to the squared-exponential kernel's intensity ω being set to 0.56. A notable trait of the covariance matrix is the high covariance values between data-points close together. The range of high covariance was determined by the squared-exponential kernel's length-scale l being set to 0.25. Comparing the uncertainty in the regression of points between the first the second data-points with that between the third and forth data-points illustrates that greater uncertainty is expected in regions of lower covariance.

A simple tool that can be utilized to assess the performance of a GP is cross-validation. Figure 5.3 demonstrates two cross-validation visual tools that can highlight potential issues within the GP. For this cross-validation single data-points are removed from the set upon which the GP is conditioned and then the GP infers values for the removed point. The comparison of the removed values with the inferred values are shown in Figure 5.3a. This comparison is excepted to be nearly linear and any points too far from the line shown with a slope of one should be considered questionable. The standardized cross-validation shown in Figure 5.3b is the cross-validated residual normalized by the cross validated standard error of the inferred point or

$$\frac{f(x^{(i)}) - \hat{f}_{-i}(x^{(i)})}{s_{-i}(x^{(i)})}, \tag{5.20}$$

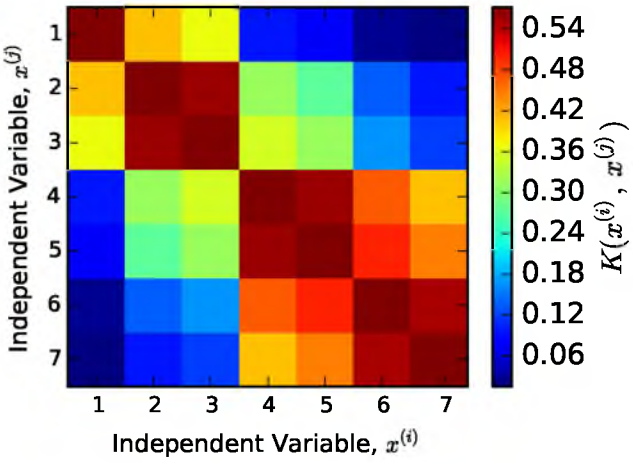

Figure 5.2: Visualization of the covariance matrix corresponding to data-points shown in Figure 5.1. Relative values of covariance between independent variable points is displayed as color quantified by the colorbar.

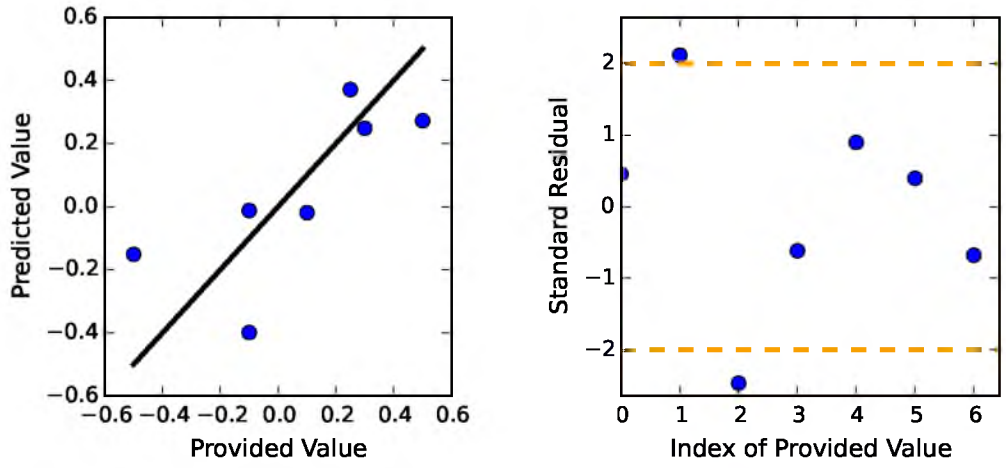

(a) (b)

Figure 5.3: Cross-validation visual tools utilized to assess GP performance. Plot (a) compares values inferred, when left out of the GP conditioning, with the true value. Black line of slope unity is provided as guideline of the desired trend. Plot (b) shows the standardized cross-validated residuals produced by leaving one data-point out of the GP and then inferring that value. Yellow dashed lines correspond to bounds for good performance behavior. Data points correspond to GP shown in Figure 5.1.

where $f(x^{(i)})$ is the data-point removed, $\hat{f}_{-i}(x^{(i)})$ is the inferred value for the removed data-point from the GP not conditioned upon the removed data-point i, and $s_{-i}(x^{(i)})$ is the standard deviation of the point inferred. According to Sacks and Welch (2010) [123], upon which these cross-validation metrics are based, the absolute standard residual typically should be less than two. Values outside that bound but within three should be of slight concern and anything larger should be flagged for further inspection.

5.1.4 Markov Chain Monte Carlo Overview

Markov chain Monte Carlo (MCMC) is a numerical technique frequently used to characterize posterior distributions when alternative means become unfeasible. MCMC can be simplistically described as a random walk whose step selection is biased by the underlying probability distribution it is tasked with exploring. That underlying probability distribution is ultimately characterized by considering the aggregate of the MCMC steps. Areas were more steps were taken indicate higher probability regions. Much more in-depth descriptions of MCMC theory and methods can be found in literature sources such as Gilks et al. (1996), MacKay (2003) and Liu (2008) [46, 88, 87].

For the present problem the MCMC Python tool emcee created by Foreman-Mackey et al. (2013) [38] will be employed. Emcee uses an affine-invariant ensemble-sampling algorithm, originally proposed by Goodman and Weare (2010) [49], to perform MCMC. The affine invariance of the sampling allows for more efficient sampling of anisotropic distributions. Ensemble sampling refers to the fact that multiple Markov chains, also known as 'walkers,' are evolved simultaneously. This parallel evolution allows the algorithm to be conducted computationally parallel, increasing the speed of sampling. Like all MCMC flavors a 'burn-in' period is necessary for the walkers to become established throughout the full distribution. This allows poor initial parameter guesses to be overcome.

Unlike the traditional Metropolis-Hastings (M-H) algorithm where step selection depends upon sampling a transitional distribution based upon the previous step, step selection for each walker is based upon the position of the complementary ensemble of other walkers [38]. Sampling from a distribution still occurs, but the sampled value is used to weight the difference between the current walker and one randomly selected from the complementary ensemble of other walkers, and it is then this weighted difference that is added to the current walker's position. This step selection method is informally known as the 'stretch move' [38]. The ratio of the probabilities of the old and new position of the walker, in a fashion similar to traditional M-H, is utilized to determine the likelihood that the new position is accepted.

Two measures of the MCMC performance are included within the emcee tool: the acceptance fraction and the autocorrelation time. The acceptance fraction is the fraction of steps proposed throughout the MCMC run that were accepted by the algorithm. Foreman-Mackey et al. [38] suggest that the acceptance fraction should be between 0.2-0.5 for an effective/efficient exploration of the distribution. The autocorrelation time is a measure of how many Markov steps are necessary to to produce independent samples. This independence measure is eased by the simultaneous evolution of multiple Markov chains. Comparing autocorrelation times with respect to each parameter in the distribution after the MCMC run allows the user to decide if further evaluations are needed to characterize the distribution. Additional evaluations can be easily gained through simply restarting the MCMC with the current walker positions.

5.2 Kennedy O'Hagan Approach

5.2.1 Kennedy O'Hagan Model-Form Uncertainty Theory

The following is a general overview of the approach presented by Kennedy and O'Hagan (2001) [74] (KOH) for tackling model-form uncertainty. Models aimed at capturing physical phenomena are typically based upon many parameters. While numerous parameters may exist in a model, the parameters that are changing within the situation of interest are those often denoted as the active parameter-set. The active parameter-set can then be split into parameters who values are known, scenario parameters x, and those whose values are unknown, calibration parameters u. Outputs of interest from models, which can be compared with observable quantities, are known as quantities of interest (QoI) and are typically represented as y.

Models are usually imperfect and often contain incomplete descriptions of the physical realities they aim to represent. The difference between the physical reality $y^R(x)$ and the model's prediction thereof $y^M(x,u)$ can be termed the model discrepancy or bias $\delta(x)$. Mathematically this relationship can be written

$$y^R(x) = y^M(x,u) + \delta(x). \tag{5.21}$$

Note that the discrepancy is only a function of the scenario parameters, not of the calibration parameters. Effectively, the 'true' discrepancy function can be thought of as corresponding to the difference between reality and the model evaluated at the 'true' calibration value.

Currently, models can only be validated if experimental data exists that can substantiate the model's basis of prediction. With experimental data there is generally some error ϵ

associated with it. Reality is usually not known, which is why experiments were undertaken in the first place, but can be described by the experiments and their associated error or

$$y^R(x) = y^E(x) + \epsilon. \tag{5.22}$$

Random/stochastic experimental error can usually be characterized through the gathering of repetitions, while capturing any bias in the experiments necessitates alternative methods. Presently, experimental bias will be ignored, while bias in the model is the focus. Following this assumption the experiment error is assumed to be independently, identically distributed (iid) and Gaussian with zero mean, but these characteristics are not a requirement for the general approach. Now removing unknowable reality from the approach by combining Eq. (5.21) and (5.22), a relationship in terms of QoI, adjusted for bias and accounting for experimental error can be validated with experimental data and utilized for predictions outside of ranges data is currently available,

$$y^E(x) = y^M(x, u) + \delta(x) + \epsilon. \tag{5.23}$$

5.2.2 Kennedy O'Hagan Algorithm

The algorithmic approach taken towards implementing Eq. (5.23) throughout this analysis entails the following steps.

- First, a GP acting as a surrogate for the imperfect model is created using a set of parameter values that adequately explore the prior parameter ranges. The hyperparameters of the covariance kernel used within the surrogate GP can either be treated with the full-Bayes method or estimated through the MAP approach,

$$P(y|X, u, hp_m) \qquad \overrightarrow{MAP} \qquad P(y|X, u, hp_{m_*}). \tag{5.24}$$

- The surrogate GP can then infer output values at input parameter positions corresponding to available experimental data or

$$P(y^*|x^*, y, X, u, hp_{m_*}) \text{ where } x^* = x_e. \tag{5.25}$$

- Any differences δ between the inferred means from the surrogate GP and the means of the experimental data can then be treated as y data inputs for the discrepancy GP. Use of the experimental mean instead of the raw experimental data is discussed in Section 5.3.6.1. The discrepancy could be modeled with alternative forms such as constant values or Gaussian random variables as was investigated by Ling et al.

(2014) [85], but GPs will be utilized throughout this example. The discrepancy GP will initially include both scenario and calibration parameters, but through analyzing the posterior distribution or potentially other characteristic distributions, calibration parameters should be defined,

$$P(\delta(x_e, u)|X_e, u, hp_\delta) \qquad \overrightarrow{calibrated} \qquad P(\delta(x_e, u)|X_e, u_\star, hp_\delta). \tag{5.26}$$

- A discrepancy GP can then be built that only depends upon scenario parameters. The hyperparameters of the covariance kernel used in the discrepancy GP can be estimated through either the full-Bayes or MAP approaches when analyzing the hyperposterior,

$$P(\delta(x_e)|X_e, u_\star, hp_{\delta_\star}). \tag{5.27}$$

- The posterior predictive distribution then allows for predictions of the QoI by adding the model and discrepancy GP's posterior predictive means and uncertainties. In order to predict full uncertainty values, the full-Bayes approaches is necessary, but this additional uncertainty is often assumed insignificant and ignored,

$$\begin{aligned} P(y^*|x^*, y, \delta(x_e), X_e, X, u_\star, hp_{m_\star}, hp_{\delta_\star}) = \\ P(y^*|x^*, y, X, u_\star, hp_{m_\star}) + P(\delta(x^*)|x^*, \delta(x_e), X_e, u_\star, hp_{\delta_\star}). \end{aligned} \tag{5.28}$$

A graphical representation of this algorithm is shown in the form of a Bayesian network within Figure 5.4. Bayesian networks are directed acyclic graphs that provide a visual representation of the underlying conditional probabilistic relationships for a Bayesian analysis. Within the Bayesian network, lines represent conditional dependence, circles are random variables, squares are data, and triangles are models. The models are not a necessary part of the network, but help provide context and ease interpretation.

5.3 Pedagogical Problem

5.3.1 General Problem Overview

The pedagogical problem is framed around modeling the reaction kinetics of

$$\mathrm{SiH_4} \rightarrow \mathrm{Si} + 2\mathrm{H_2}. \tag{5.29}$$

To predict the progression of this reaction the following equation can be solved

$$\frac{\mathrm{d}y_{\mathrm{SiH_4}}(t)}{\mathrm{d}t} = -u\ y_{\mathrm{SiH_4}}(t), \tag{5.30}$$

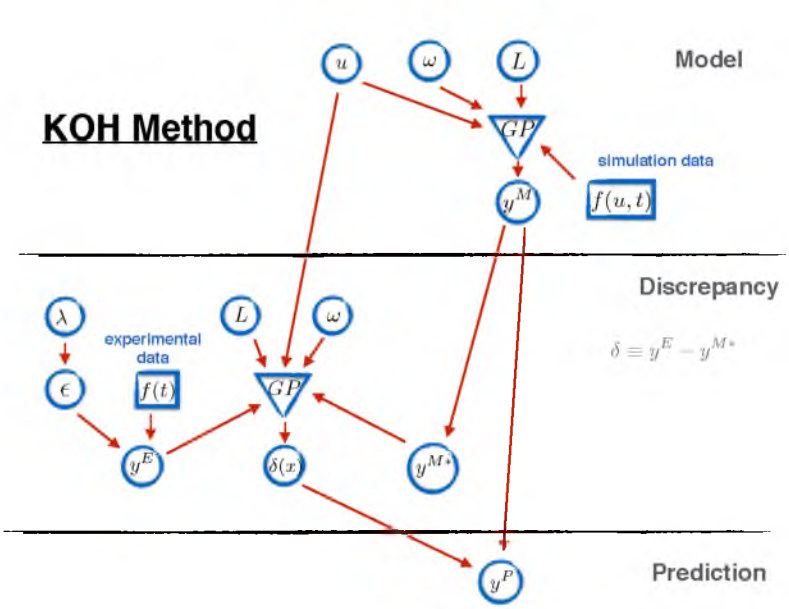

Figure 5.4: Bayesian network representation of algorithm used to implement the KOH approach to model-form uncertainty and calibration. Here circles represent random variables, squares correspond to data, triangles indicate models, and lines represent conditional dependence.

where $y_{\mathrm{SiH}_4}(t)$ is the concentration of SiH_4 at time t, u is an unknown reaction-rate and the initial condition is $y_{\mathrm{SiH}_4}(0) = y_{0,\mathrm{SiH}_4}$. The solution to this equation is

$$y_{\mathrm{SiH}_4}(t) = y_{0,\mathrm{SiH}_4} \exp(-u\, t), \tag{5.31}$$

but this is known to be an incorrect model-form. The actual equation to describe the aforementioned kinetics is

$$y_{\mathrm{SiH}_4}(t) = (y_{0,\mathrm{SiH}_4} - c) \exp(-u\, t) + c, \tag{5.32}$$

where c is the concentration of SiH_4 remaining at the end of the reaction (c = 1.5), also known as the residual.

Although Eq. (5.31) is known to contain model-form error, this exercise will explore how well the model with incorrect physics can be utilized to make predictions after implementing the model-form uncertainty validation scheme described within Kennedy and O'Hagan (2001) [74]. Following this approach it can be seen a priori that a model discrepancy or deficiency term would analytically be described as

$$\delta(t) = c\,(1 - \exp(u_\diamond\, t)), \tag{5.33}$$

where $u_\diamond$ designates the true value of the reaction-rate being calibrated and whose value is 1.7.

Parameters explored throughout this exercise are time t and the reaction-rate parameter u, which are treated as scenario and calibration parameters, respectively. The prior parameter ranges are $[0-3]$ for both parameters. Such bounds could also be thought of as uniform prior distributions. Utilizing the points selected by Bayarri et al. (2005) [11], which were determined by a Latin hypercube design of experiments (DOE), Eq. (5.31) can be evaluated at those points, providing output $y_{\mathrm{SiH}_4}(t, u)$. The parameter space explored by this DOE was [0.5, 2.0] x [0.1, 3.0] for u and t, respectively, and the design points are shown in Table. 5.1. The initial condition used throughout the example is $y_{0,\mathrm{SiH}_4} = 5.0$. Latin hypercubes are a popular method of efficiently exploring parameter spaces with high dimensionality, while also being computationally affordable. Compared to randomly sampling the parameter domain, Latin hypercubes have smaller variance on their sample mean [31].

5.3.2 Application Approach

The parameter design points and corresponding model outputs can be used to construct a GP surrogate model utilizing the squared-exponential covariance kernel. Although evaluations of Eq. (5.31) are computationally cheap, a surrogate model will still be utilized to demonstrate how this analysis could be conducted for models with expensive evaluations. Optimal values for the length, l_1 and l_2, and amplitude, ω, hyperparameters can be approximated using the MAP method. Log-normal prior distributions with means 0.5, 0.7, 0.2 and standard deviations 0.2, all reported in normal space equivalent, are specified for hyperparameters ω, l_1, and l_2, respectively. No noise kernel is included in this GP due to the model providing deterministic outputs. This GP surrogate along with the model it is based upon are compared against the true physics in Figure 5.5. The model outputs corresponding to samples spanning the full prior parameter range are shown within Figure 5.5a and the best case scenario where the true calibration value 1.7 was used is shown in Figure 5.5b. It can be noted that the GP performs well in regions were data was sampled, but fails to

Table 5.1: Latin hypercube design of experiments' selected parameter points (time and reaction-rate u).

Latin Hypercube Design of Experiments										
Time	2.159	0.941	0.303	0.709	1.753	1.144	0.506	2.391	1.956	1.550
	2.594	2.797	1.347	0.100	3.000					
u	1.145	2.000	0.710	1.040	1.895	0.605	1.685	1.565	0.500	0.935
	0.815	1.790	1.460	1.355	1.250					

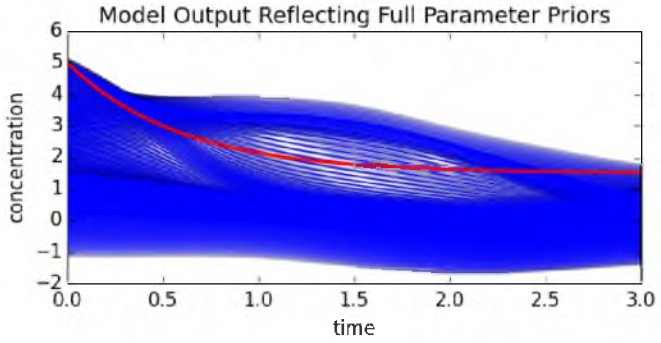

(a)

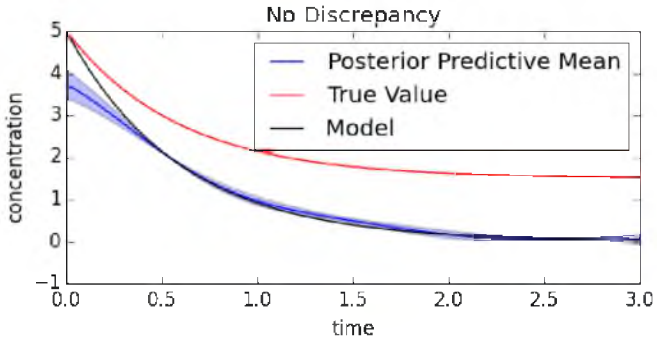

(b)

Figure 5.5: Comparison of the 'truth' (red line) with the model containing an incorrect model-form (black line) and the posterior predictive mean and one standard deviation (blue line/region) of the GP surrogate model. Plot (a) shows the possible model outputs reflecting samples spanning the full range of parameter prior (0-3). Plot (b) depicts the best case scenario where the true calibration parameter value 1.7 is used in the model and surrogate GP.

capture the initial temporal behavior due to data not being sampled in that region. The discrepancy between the model and true physics is apparent.

Synthetic experimental data were created by sampling the true physics model and adding Gaussian white-noise with a standard deviation of 0.3, Eq. (5.32) + $\mathcal{N}(0, 0.3)$. These data include three replicates for each position sampled from the parameter space and are shown within Figure 5.6 and Table 5.2. In the spirit of assuming the data were of experimental origin, the variance of the plotted experimental data was calculated following the assumption that the data were normally distributed and the mean and one standard deviation are also shown with the data-points. This same assumption was utilized in modeling the experimental noise within the noise kernels throughout the exercise as previously mentioned.

A GP of the discrepancy between the model GP and the experimental data can now be constructed. After inferring values with the model GP for the temporal locations of

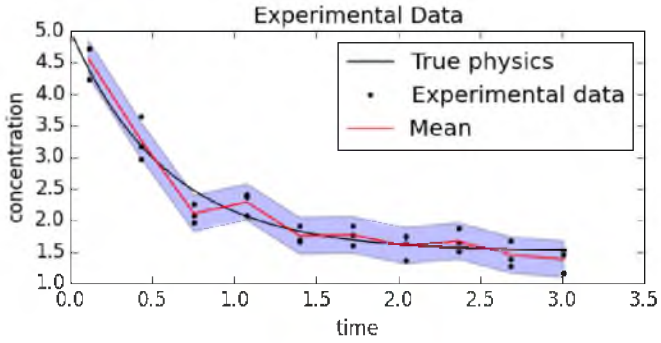

Figure 5.6: Experimental data are shown (black dots), along with its respective mean (red line) and one standard deviation following the assumption of a normal distribution (blue region).

Table 5.2: Experimental data times and corresponding three repetitions of SiH_4 concentrations.

	SiH_4 Concentration									
Time	0.110	0.432	0.754	1.077	1.399	1.721	2.043	2.366	2.688	3.010
y_1	4.730	3.177	1.970	2.079	1.908	1.773	1.370	1.868	1.390	1.461
y_2	4.720	2.966	2.267	2.409	1.665	1.603	1.661	1.505	1.275	1.157
y_3	4.234	3.653	2.084	2.371	1.685	1.922	1.757	1.638	1.679	1.530

the experimental data, those values can be compared with the experimental mean values and the difference utilized as y inputs to create a discrepancy GP. Figure 5.7 depicts a range possible discrepancy GPs, each corresponding to a different calibration value, which is still unknown. For each discrepancy GP shown, a noise kernel and squared-exponential kernel with two length-scales were used and their respective hyperparameters estimated via the MAP method. When calculating the posterior, the prior for the noise intensity is set to the maximum standard deviation of the experimental data and the priors for squared-exponential's intensity and length are both set to log-normal with mean/standard deviation 0.5/0.2 and 0.9/0.2, respectively, in normal space equivalent.

5.3.3 Calibration Estimation Techniques

The discrepancy GP can only provide help with extrapolative predictions if the true calibration parameter value is utilized when constructing the discrepancy GP. If the discrepancy is based upon the wrong calibration parameter value, it will only provide useful predictions near the validation data parameter space. The ideal manner of searching for the calibration parameter is to search through the full posterior distribution. This could be accomplished with an MCMC algorithm, but such a search is typically computationally costly. A few alternative approaches for approximating this search also exist.

The first calibration estimation method considered is based upon the discrepancy GP's hyperposterior once the hyperparameters have been estimated and is shown in Figure 5.8. This method is convenient because the calculation is already performed when utilizing the MAP hyperparameter estimation technique. Using estimated hyperparameter values allows for a psuedo-marginalized hyperposterior to be explored over the calibration parameter range,

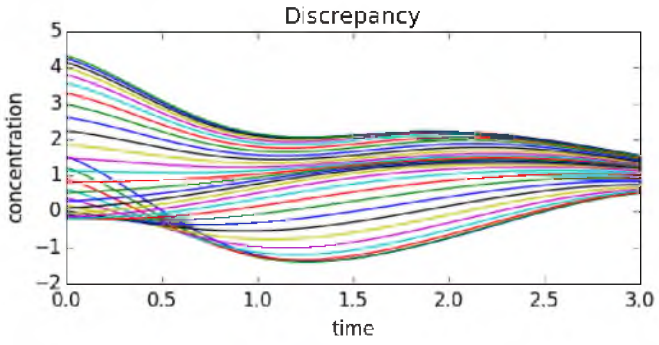

Figure 5.7: Possible discrepancy functions where each line reflects a different value for the unknown calibration parameter between 0-3.

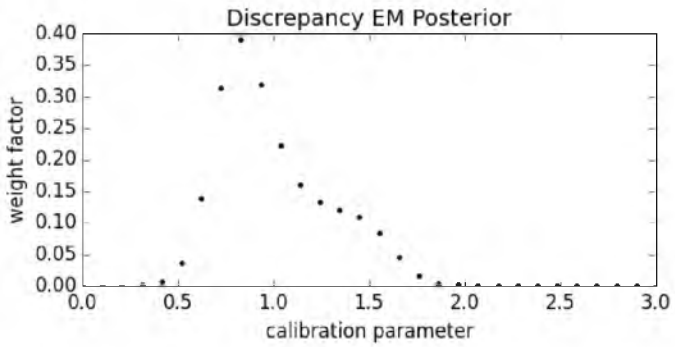

Figure 5.8: The discrepancy maximum hyperposterior over a range of calibration values after the hyperparameters were estimated with MAP. The weight factor is effectively an unnormalized probability.

$$P(u|y, X, K(X, X), hp_{\text{MAP}}) = P(y|X, K(X, X), u, hp_{\text{MAP}})\ P(u). \tag{5.34}$$

The hyperposterior is referred to as psuedo-marginalized here because the MAP algorithm is performed for each discrete calibration value explored. This method is generally known as the expectation-maximization (EM) algorithm. Figure 5.8 shows the distribution produced by this method and using an averaging based upon distribution weight, a calibration value can be estimated. This method resulted in a calibration value estimate of 0.997. As is generally the case with most distributions found, the use of the weighed mean value was only one of the possible calibration estimates that could be chosen from this distribution. For the MAP calculations the noise was set equal to the maximum standard deviation in the experimental data and the squared-exponential's intensity and the squared-exponential's intensity and length were both set to log-normal with mean/standard deviation 0.5/0.2 and 0.9/0.2, respectively, in normal space equivalent.

Another means of estimating the calibration parameter is to find the minimum sum of squared errors for the discrepancy function over the calibration range. This can be seen as approximation of the integral of each of the possible discrepancy functions in order to locate the smallest

$$u_* \rightarrow \min\big(\sum_i \delta(x_{e,i}, u)\big). \tag{5.35}$$

This method is shown within Figure 5.9a and results in a calibration estimate of 0.72. This method is not deemed the best approximation due to there being no basis for why it should be assumed that the smallest amount of discrepancy would be preferable. The real goal when approaching model-form uncertainty problems is to locate a discrepancy that corrects for a model based upon the true values of the calibrated parameters. Interestingly, the

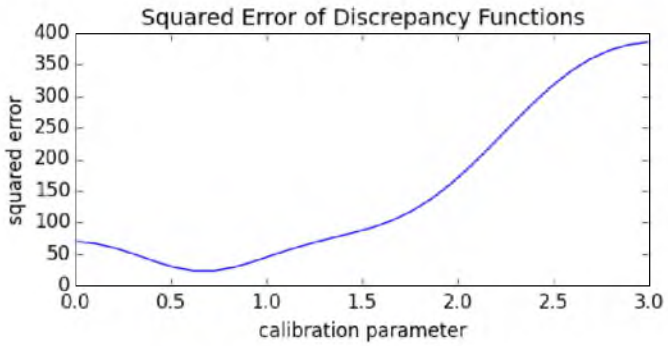

(a)

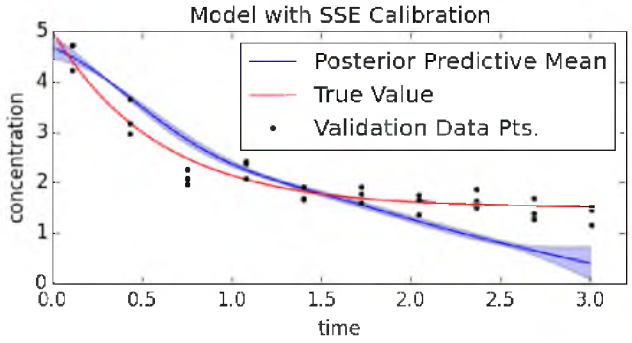

(b)

Figure 5.9: The sum of the squared-errors (SSE) of discrepancy functions for a range of calibration values is shown in plot (a). Plot (b) shows the corresponding GP mean and one standard deviation (blue line/region) based upon the calibration value 0.72 found with the SSE method versus the true physics (red line) and validation data-points (black dots).

minimized discrepancy calibration estimate should be similar to how parameter calibration of the model would proceed if a discrepancy term was not utilized, also known as a residual minimization method. The only significant difference being that noise was included in the discrepancy functions. It can be seen within Figure 5.9b that the model calibrated without bias would perform poorly even in a regressive capacity.

One other possible means of estimating the calibration parameter, which in this instance can be seen to be similar to the EM algorithm estimate, is to calculate a marginal likelihood distribution where the calibration value is not marginalized over. Figure 5.10 shows this marginal likelihood distribution where the noise kernel's λ and the ω of the squared-exponential kernel have both been marginalized over, or treated as nuisance parameters

$$P(y|u, X, K(X,X), L_{\text{fixed}}) = \iint P(y|u, X, K(X,X), \omega, \lambda, L_{\text{fixed}})\, \mathrm{d}\omega\, \mathrm{d}\lambda. \tag{5.36}$$

For this marginalization, the hyperparameter length-scale was set to a fixed value of 0.9 that was hypothesized as being reasonable from an initial analysis of the data. This length-

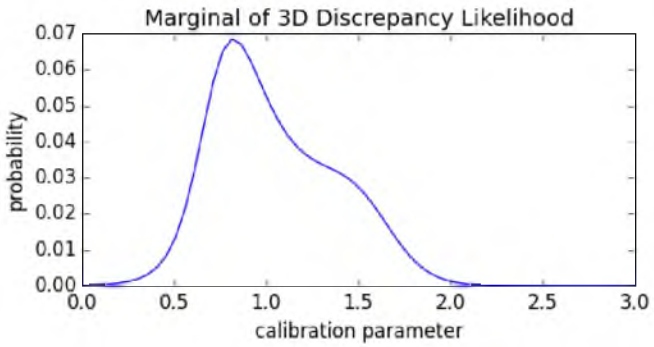

Figure 5.10: The likelihood distribution when marginalized over λ of the noise and ω of the squared-exponential kernel. The length-scale was fixed at the value of 0.9 that was estimated after considering the available data.

scale estimate was reconsidered in further analysis discussed in Sections. 5.3.5.2-5.3.5.3. The fact that this distribution strongly resembles the psuedo-marginalized hyperposterior distribution found with the EM approach, as seen by comparing Figure 5.8 and 5.10, is a positive for the EM algorithm approach, due to it being computationally cheaper to compute. Here the marginalization method estimates the calibration parameter to be 1.04. Weight averaging of the distribution was again used.

5.3.4 Prediction

Once the calibration parameter has been estimated, it can be utilized to create a discrepancy GP dependent only upon scenario parameters $\delta(x)$, which when added to the original model GP allows for predictions as shown in Figure 5.11. Compared to the original model predictions, Figure 5.5, the model adjusted by accounting for the discrepancy has superior performance within the temporal range covered by the validation data. Even where the mean of the posterior predictive does not directly line up with the true model, 'truth' is within a standard deviation.

5.3.5 Continued Analysis

While the results shown in Figure 5.11 appear promising, recall that the calibrated reaction-rate parameter value used, as found in Figure 5.8, was 0.997, while the actually value is known to be 1.7. This means that the model shown in Figure 5.11 will likely be inaccurate outside of the range of the validation data. This conclusion can already be seen to an extent in Figure 5.11 near the end of the temporal range where the posterior predictive mean begins to drift downwards and has a negative slope. Physically it is known that the system is limited by a chemical equilibrium that maintains a residual amount of

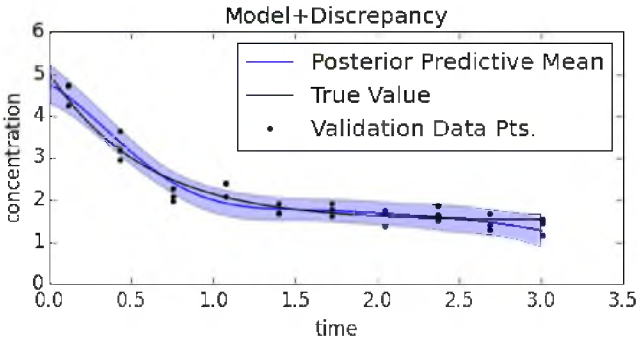

Figure 5.11: The mean and one standard deviation of the posterior predictive distribution generated by adding the model and discrepancy GPs (blue line/region), both based upon the calibration parameter value (0.997) estimated by the EM approach. The data used for validation (black dots) and the true model-form (black line) are also shown for comparison.

SiH_4 (c = 1.5), but the GP surrogate models are not based upon physical constraints. Such physical constraints could be built into the GP through a prior mean, but that would imply prior knowledge about the system not currently assumed. Filtering random samples of this posterior predictive distribution by placing a set of constraints upon them could be another means of improving predictive capabilities (Section 5.5.3).

5.3.5.1 Model DOE

An issue that should be addressed is the fact that the DOE of the model inputs did not cover the full temporal range, but started at 0.1. Considering that such DOE is in the hands of the computational modeler and not limited by available experimental data, it seems reasonable to ensure that the full parameter prior distribution is explored. To investigate any effects this would have upon the results, the design point (0.100, 1.355) was switched to (0.000, 1.355). This alternation was found to have a minor effect upon the initial times of the final predictive distribution, which generally appears to be more strongly effected by the validation data. There were small changes in the calibration estimates with the EM estimate changing from 0.997 to 1.1 and the marginal distribution estimate changing from 1.04 to 1.09.

5.3.5.2 Three-Dimensional Likelihood

A grid-based evaluation method was used to generate a three-dimensional likelihood distribution. Two-dimensional marginals and a slice of the full three-dimensional distribution are shown in Figure 5.12. All marginalization was undertaken by using the cumulative trapezoid method to integrate over dimensions. The three-dimensional likelihood that these

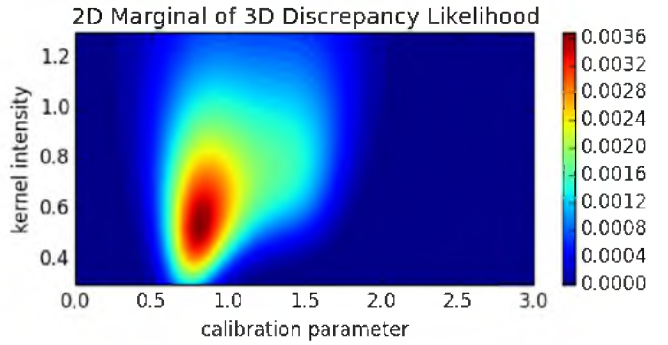

(a)

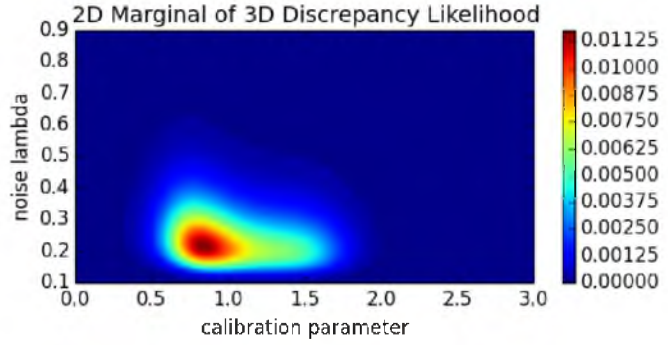

(b)

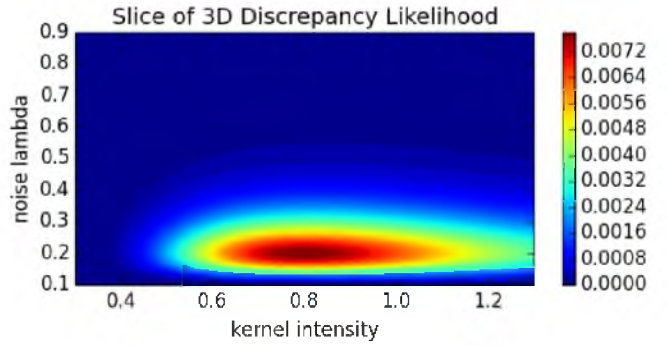

(c)

Figure 5.12: Two-dimensional marginal likelihood distributions (a, b) and a slice of the three-dimensional discrepancy likelihood (c) that the marginals are based upon. The likelihood was marginalized over λ of the noise kernel for the plot (a), over ω of the squared-exponential kernel of the plot (b), and was sliced through the calibration dimension at the median parameter range value (1.5) in plot (c). The length-scale was fixed at a value of 0.9.

plots are based upon was parameterized by the calibration parameter u, λ of the noise kernel, and ω of the squared-exponential kernel. The length-scale was set at the mean prior value 0.9 because this length-scale was initially deemed to be a reasonable estimate after considering the data spacing.

From these two-dimensional marginals and slice of the three-dimensional likelihood, the calibration value estimated by the EM method appears to be reasonable and there is no indication within these distributions that would lead to any inference of the true calibration value. The marginal displaying the distribution over the squared-exponential kernel intensity and the calibration parameter indicates some correlation by its slight diagonal skew. There is minimal and no evidence of correlation between parameters for the other marginal and within the slice. Any correlation noted between the parameters in a posterior equivalent of Figure 5.12 would not conflict any assumptions regarding the independence of the the parameter's prior distributions because such correlations were only learned through the Bayesian process. Due to the low probability found in the region of the known true calibration parameter, the assumption of using a fixed length-scale to create the likelihood was reconsidered.

5.3.5.3 Four-Dimensional Likelihood

A four-dimensional likelihood including the same parameters considered within the previous three-dimensional likelihood with the addition of the squared-exponential length-scale is now considered. Two-dimensional marginal distributions, shown in Figure 5.13, display many attributes not contained in equivalent marginals based upon a three-dimensional likelihood. Regions of high probability found in Figure 5.12, now appear to be local phenomena, but not the true global description. The noise variable λ still appears to be independent of the other parameters. The length parameter l shows clear correlation with the calibration parameter, while the squared-exponential intensity appears to have a weaker or minimal dependence. Perhaps the most interesting information to be gained is within the marginal distribution shown in Figure 5.13d, where all three hyperparameters have been integrated over, leaving a bimodal distribution. This bimodal distribution shows the region found in Figure 5.10 to be a local maximum and that a global maximum also exists with a value of 1.69, or approximately the true value. This behavior is also present in all three corresponding two-dimensional marginals. Unlike previous estimation of the calibration parameter where weighted means were used, a maximum value appears more appropriate here considering the distribution is bimodal.

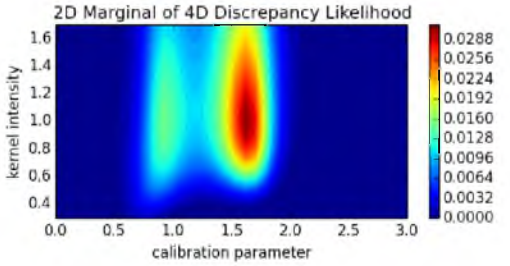

(a)

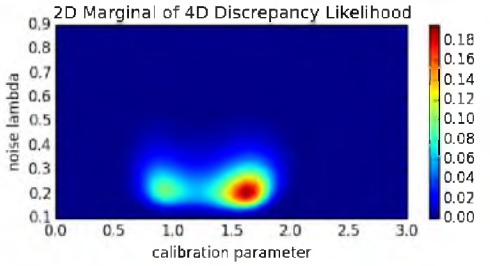

(b)

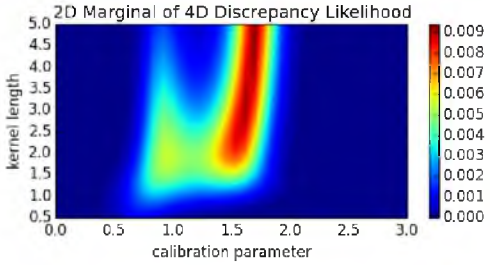

(c)

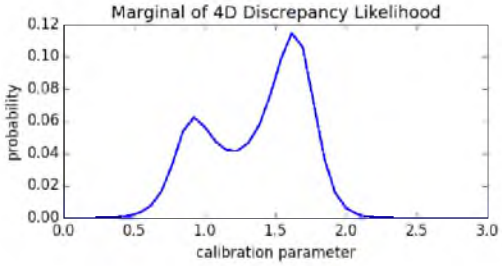

(d)

Figure 5.13: Two- and one-dimensional marginals of the discrepancy's four-dimensional likelihood distribution. Plot (a) shows the distribution of the calibration parameter and the squared-exponential kernel's intensity ω. Plot (b) show the distribution for the calibration parameter and the noise parameter λ. Plot (c) plot shows distribution for the calibration parameter with the squared-exponential's length hyperparameter l. Plot (d) shows the one-dimensional marginal, where the three hyperparameters, λ, ω, and l have been integrated out leaving only the unnormalized probabilities across calibration parameter space.

5.3.5.4 Analytic Discrepancy Comparison

Another comparison that exposes potential issues with making extrapolative predictions is between the discrepancy GPs created and the known analytic form as shown in Figure 5.14. Here it can be seen that when using the calibration value found with the EM algorithm method (0.997) that the discrepancy GP has different features than the analytical solution, Figure 5.14a. Even when using the known calibration value (1.7) to create the discrepancy GP, it can still be seen in Figure 5.14b that the method behind creating the discrepancy term does not yield the correct form, but does have values closer to the analytic solution throughout most of the temporal parameter range. From these plots it can be concluded that the model with discrepancy included would be unlikely to be predictive outside of the parameter range of the validation data if the correct calibration value was not found and even if the correct calibration parameter value is found, not a great deal of confidence should be placed upon its extrapolation.

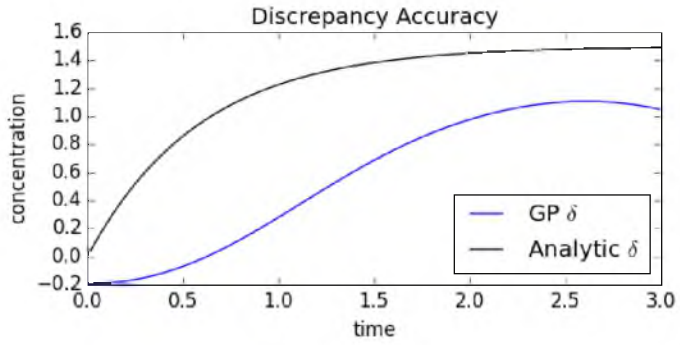

(a)

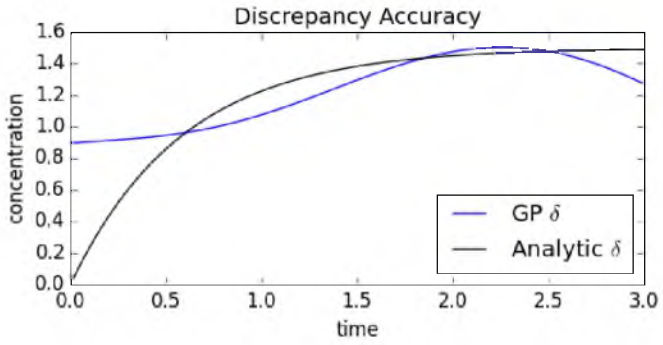

(b)

Figure 5.14: A comparison of the discrepancy GP mean (blue line) with the analytical discrepancy function (black line). In plot (a) the discrepancy GP uses the calibrated parameter found with the EM algorithm method (0.997) and plot (b) uses the true value of the calibration parameter (1.7).

That being said, both discrepancy models do provide insight into the shortcomings of the original model containing incorrect physics. Both discrepancy GPs indicate that the model was under-predicting the concentration and that qualitatively this increased over time until the end of the time domain. Looking back to the original model Eq. (5.31), there appear to be three potential locations within the model that the discrepancy could be attributed to: acting linearly along with y_0, acting exponentially with u and t or as an additional term added on. Clearly the discrepancy has temporal functionality, and at least a portion must be additive to the current model-form in order to level out at a value other than zero. Hypotheses about model-form error within the preexponential and exponential terms would require further consideration.

5.3.5.5 Posteriors and Priors

Following the Bayesian philosophy, studying the discrepancy's posterior distribution is the next logical step. In comparing the posterior and likelihood distributions the effect of the priors can be seen. Initially, the discrepancy's posterior was created using a grid-wise evaluation of the parameter space, just as was previously done with the likelihoods except that hyperparameter priors were now included. The effect of slight alterations to the priors can be noted in posterior marginals shown in Figure 5.15. Here it can be seen that by changing the standard deviation and mean parameters of the log-normal distributions used for the hyperparameter priors of the squared-exponential kernel causes significant changes to the ultimate posterior marginals found for the calibration parameter. All log-normal distribution parameters are described in their normal space equivalent due to that space generally being more intuitive. Simply changing the standard deviation between 0.1, 0.2, and 0.5 (plot (b), (a), and (d)) causes two modes of the distribution to become differentiated. Changing the mean from 0.9 to 0.5 (plot (a) and (c)) can also be seen to have a significant effect upon the presence of two modes within the distribution. The shift of 0.4 to the mean of the prior distributions effectively caused the second mode of the original likelihood distribution to be ignored.

Analyzing posterior distributions in many dimensions becomes computationally expensive, often leading researchers to utilize MCMC methods. Following this tradition, the MCMC tool of Foreman-Mackey et al. (2013) [38] was then utilized to further explore the posterior distribution. Upon implementation the MCMC results were verified against the grid-wise results (Figure 5.16). In order to avoid allowing the experimental data to effect the surrogate model GP's characteristics, a modular approach towards probabilistic exploration

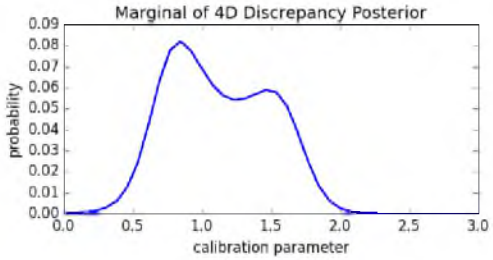

(a)

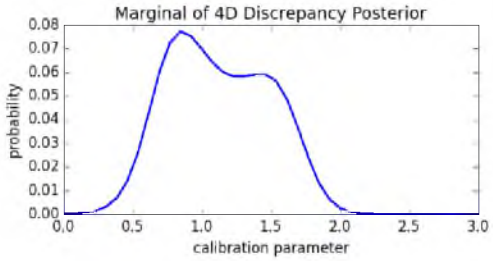

(b)

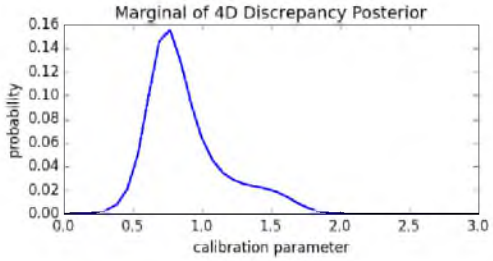

(c)

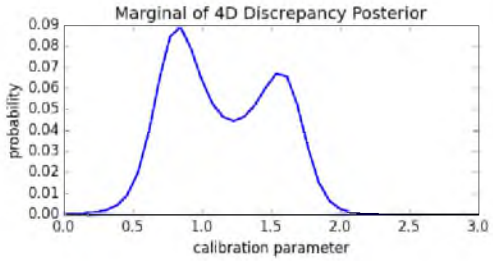

(d)

Figure 5.15: Demonstrating the effect of prior distributions on the posterior distribution. For these four plots the same likelihood used to generate Figure 5.13 was implemented, but with a Jeffreys prior on the noise intensity λ and log-normal distributions for the squared-exponential kernel's intensity ω and length l. The standard deviation of the squared-exponential's hyperparameter priors, described in terms of an equivalent normal distribution, was altered between plots to values 0.2, 0.1, 0.5, and 0.2 for plots (a), (b), (d), and (c), respectively. The mean of those priors was 0.9 for all plots except plot (c), which was 0.5.

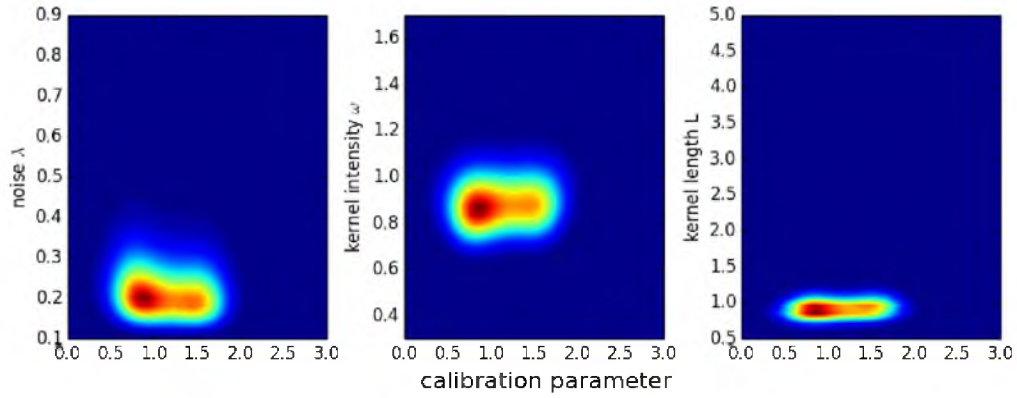

(a)

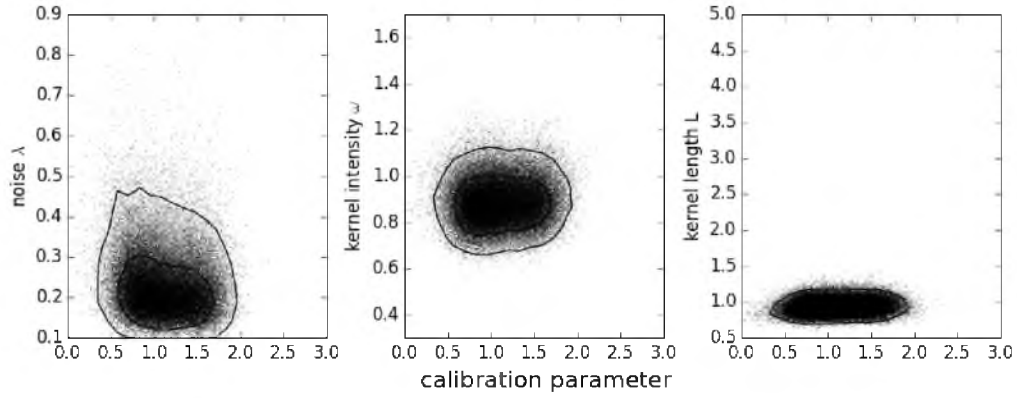

(b)

Figure 5.16: Demonstrating that the MCMC run using the emcee tool from Foreman-Mackey et al. (2013) [38], plot (b), is able to produce results equivalent to those generated via grid-wise search, plot (a). The compared posterior marginals used the same hyperparameter configuration as Figure 5.15b.

should be taken [11]. Such modularity refers to separately exploring the surrogate and discrepancy GPs parameter spaces, and the following MCMC explorations only involved the discrepancy.

Once the MCMC method was verified, different configurations of the hyperparameter priors were analyzed. These MCMC runs all utilized 100 walkers or separate Markov chains, with a 1,000 step burn-in period and a total of 3,000 Markov steps. These MCMC runs generally had a mean step acceptance ratio around 0.4. Figure 5.17 demonstrates that small changes to the prior distributions can have significant effects upon the marginal posterior spaces.

As would be expected, the bottom plot of Figure 5.17 shows that the Jeffreys prior, which is meant to be an uninformative prior, produces marginals most similar to the likelihood

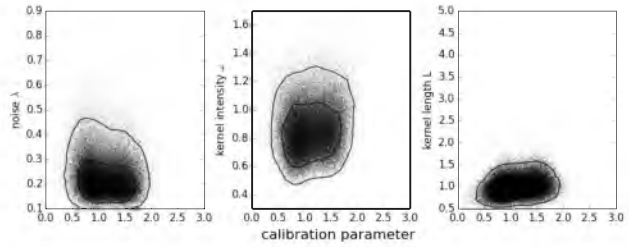

(a)

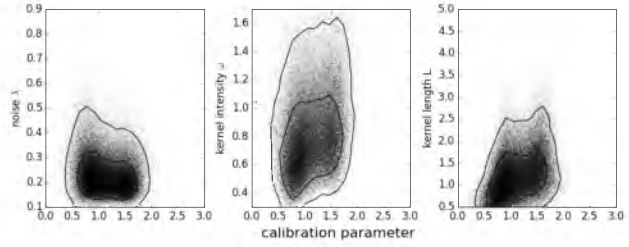

(b)

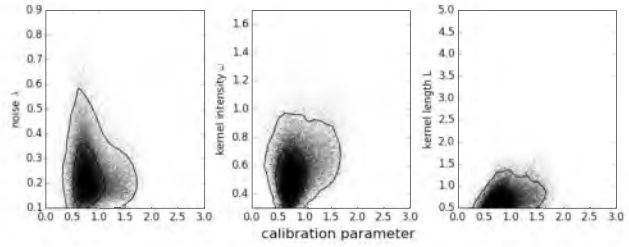

(c)

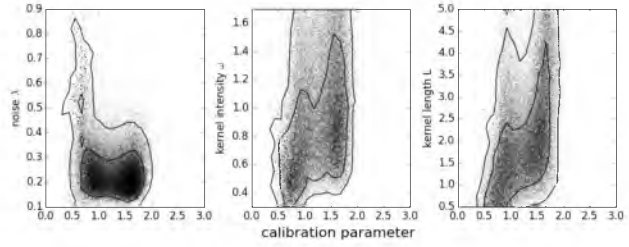

(d)

Figure 5.17: Comparison of marginals of the discrepancy posterior with varied priors on the hyperparameters. All plots included a Jeffreys distribution for the λ in the noise kernel. Plots (a), (b) and (c), plus Figure 5.16b, used log-normal distributions for the squared-exponential's intensity ω and length l priors. All log-normal parameters are described in their normal space equivalent. Plot (a) had a mean of 0.9 and std of 0.2, plot (b) a mean of 0.9 and std of 0.5, plot (c) had a mean of 0.5 and std of 0.2, Figure 5.16b had a mean of 0.9 and std of 0.1, and plot (d) used Jeffreys distribution for both ω and l.

equivalents previously shown in Figure 5.13. The log normal distribution produces marginals least similar to the likelihood equivalents when centered at 0.5. Being a more informative prior, the log normal distribution is showing a strong effect upon the posteriors. Thus, the marginals corresponding to log normal priors centered furthest from the second mode in the likelihood distribution appear least like their likelihood equivalents. Within the three sets of marginals where the log normal priors' standard deviations were varied, the variance in the marginals directly reflects those variances. Overall, these examples demonstrate that an informative prior can have a strong effect upon the posterior distribution. As long as the prior information is correctly utilized this should ease locating a posterior, but care must be taken.

5.3.6 Initial Conclusions

Although much of the previous analysis would not be possible for real application problems, where 'truth' is not known, it does foreshadow an additional problem that could play an important role in any application of the model-form uncertainty method utilized thus far, the identification problem. The identification problem is the lack of ability to separately calibrate model parameters and create a functional model discrepancy. In the present application, the correct calibration value was found, but with uncertainty, through analyzing the discrepancy's full likelihood distribution or posterior depending upon priors selected (including calibration parameters and hyperparameters), but not with any of the alternative approximations presented. Although no known method of solving the identification problem have been presented, discerning if an identification problem exists can be studied and discussed [6]. This will be elaborated upon within Section 5.4.

The final posterior predictive distribution is shown in Figure 5.18, where the calibration value (1.69) found from the four-dimensional likelihood distribution was utilized along with hyperparameter values then estimated with the MAP method. Compared to the posterior predictive shown in Figure 5.11, the mean in Figure 5.18 appears more likely to be predictive outside of the validation data's parameter range because it is not trending away from the true physics at the end of the temporal region and generally contains more physical monotonicity.

5.3.6.1 Experimental Data Considerations

When calculating the discrepancy between the surrogate model's mean output and the experimental data, the mean value of the experimental data repetitions has been utilized throughout this exploration. The experimental data could have been used directly and

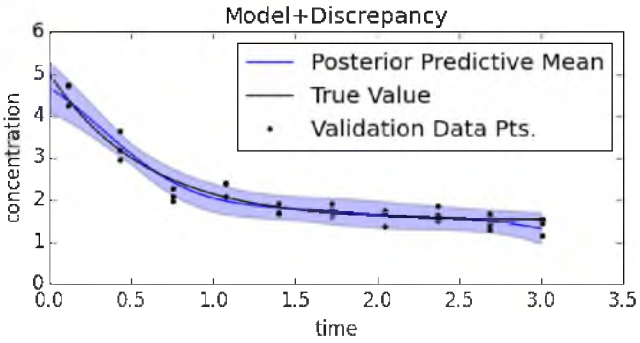

Figure 5.18: The mean and one standard deviation of the final posterior predictive distribution generated by adding the model and discrepancy GPs (blue line/region), both based upon the optimal calibration parameter value (1.69) found as the maximum of the discrepancy's four-dimensional likelihood distribution. The data used for validation (black dots) and the true model-form (black line) are also shown for comparison.

would have led to a more typical interpretation of the noise variance found within the discrepancy GP. When comparing the surrogate mean with the experimental mean the noise represents noise in moment space. If comparing directly with the experimental data the noise represent experimental data noise, or when assuming Gaussian white-noise it characterizes the variance in the distribution of repetitions. While characterizing noise in the experimental data is typically a desired goal, using the noise in the means is not believed to alter other results found throughout this analysis.

5.4 Identification Problem / Sensitivity Analysis

The identification problem is the inability to identify unique calibration parameter values while also defining model discrepancy. Calibration parameters represent underlying physical characteristics whose values contain uncertainty, thus they can be treated as random variables. Calibration of physical parameters in the context of Bayesian analysis is known as the inverse problem and is an important scientific endeavor in of itself.

5.4.1 Gradient Analysis

To tackle the identification problem the gradients of the discrepancy GP will be considered. Sharp gradients in the discrepancy GP with respect to the calibration parameter should imply that the discrepancy function and the calibration parameter are separately identifiable, while the opposite implies that identification is unlikely. This is not a means of identification, only a means of exploring if identification would be probable for the problem of interest. Following the order of analysis undertaken thus far, the model and discrepancy

GPs based upon the EM algorithm calibration estimation will be analyzed first. This same analysis based upon the four-dimensional likelihood calibration estimation will subsequently be undertaken.

Figure 5.19 shows the gradients of the discrepancy GP with respect to time and the calibrated reaction-rate parameter. The discrepancy GP is the same as that used to create the results for Figure 5.11. The gradients with respect to the calibration parameter are those of interest for the identification problem, but the gradients with respect to time are included as a means of providing relative scale for comparison. The gradient of interest has a maximum value of similar order to the equivalent gradient with respect to time. This does not appear to give credence towards any conclusions about identification. The Hessian is also included in Figure 5.20 to look for further correlation between the parameters.

The gradients of the GP acting as a surrogate model for the model containing incomplete

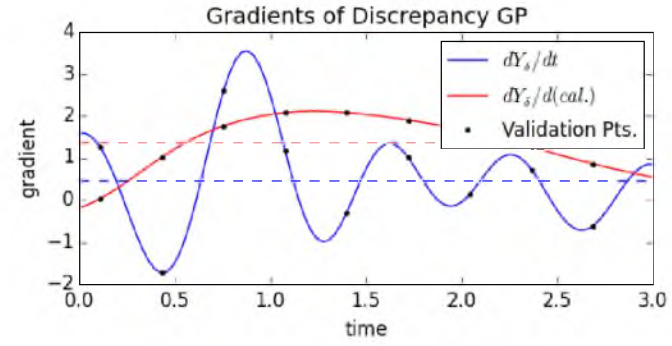

(a)

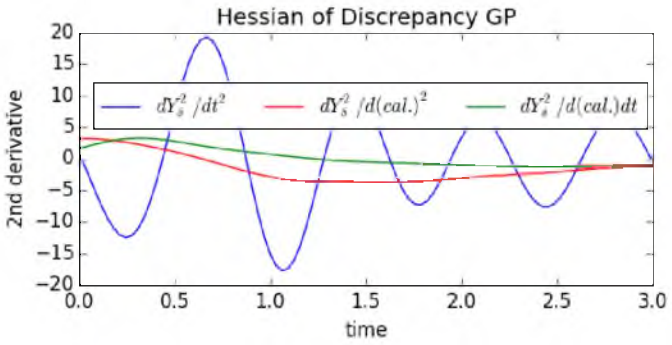

(b)

Figure 5.19: Gradients (a) and Hessian (b) of the discrepancy function along time. Dashed lines denote averaged values, blue lines are with respect to time, red lines are with respect to the calibration parameter, the green line is a mixed differentials with respect to time and the calibration parameter, and the black dots show the experimental data-points. The calibrated parameter value utilized (0.997) was found with the EM method and hyperparameters were solved for with the MAP method.

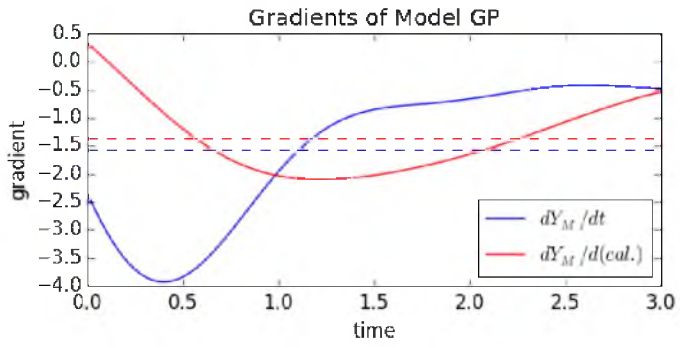

Figure 5.20: Gradients of the model GP along time. Red lines denote gradients with respect to the calibration parameter, blue lines are with respect to time, and the dashed lines are averages. The calibrated parameter value utilized (0.997) was found with the EM method and the hyperparameters were found with the MAP method

physics (Figure 5.20) can also be studied for parameter sensitivity and to look for correlation with those in the discrepancy GP. The inflection of the gradient with respect to the calibration parameter for the model is opposite in sign and of similar absolute maximum magnitude compared to the equivalent for the discrepancy function. While abstractly analyzing the model gradients for this applications bears little fruit, it could prove helpful in applications with higher dimensionality of unknowns.

While it was hard to interpret the likelihood of identification when studying the gradients of the GPs based upon the wrong calibration value, the counterparts based closer to the true calibration value (1.69 from the four-dimensional likelihood) are shown in Figure 5.21. The discrepancy GP gradients with respect to the calibration parameter have greater amplitude and magnitude when the GPs are based upon the correct calibration value. Gradients of the discrepancy GP with respect to the calibration parameter have the largest absolute magnitudes at the initial times. This could be interpreted to infer that it is these initial temporal periods where identification is most likely due to the function being most sensitive to the calibration parameter values within this temporal region. Although it has been previously mentioned that these results do not provide a definitive answer to the identification problem, it does appear that when the model was properly calibrated the gradients have greater magnitude and amplitude and thus identification appears more likely.

One final step in the examination of the discrepancy and model gradients with respect to the calibration parameter is to visualize how they change over the full calibration parameter space (Figure 5.22). For both the model and discrepancy GPs, the surfaces representing the gradients with respect to the calibration parameter over the full parameter ranges are

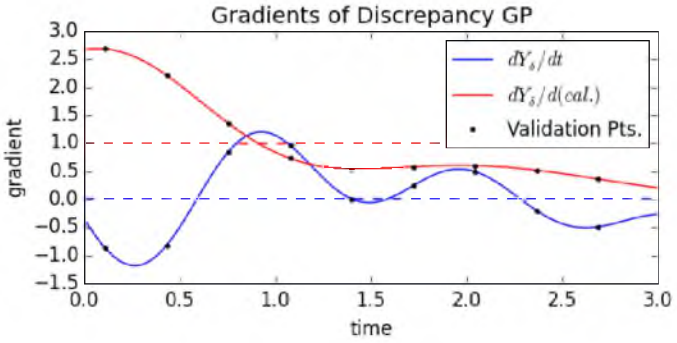

(a)

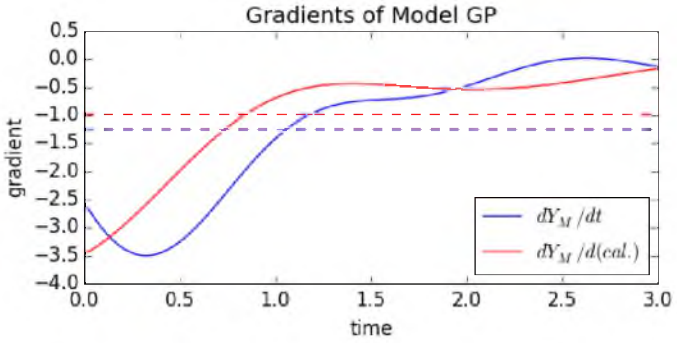

(b)

Figure 5.21: Gradients along time for the model (a) and discrepancy (b) with respect to the calibration parameter and time. Both GPs are based upon the calibration value found with the four-dimensional likelihood (1.69). Hyperparameters were located with the MAP algorithm.

near their minimal values and located in the relatively flattest regions for the calibration parameter value found through the EM algorithm approach. The true calibration value lies in an area of the surface with greater gradient values. Further extensions of this analysis are needed to gain deeper insight into this type of analysis' value towards identification issues.

5.4.2 Types of Nonidentifiability

Two categories of nonidentifiability have been proposed in the literature: structural and practical [116]. Structural nonidentifiability is caused by redundancy in the parameterization of the of the model structure [85]. Practical nonidentifiability is due to insufficient quantity or quality of experimental data. While structural nonidentifiability will not be investigated, practical nonidentifiability can easily be considered through varying the amount of experimental data as well as the aleatoric uncertainty in those data.

Figure 5.23 demonstrates the effect of data quantity. Here it can be seen that by removing three data-points (a), the distribution appears to diverge into two separate distributions,

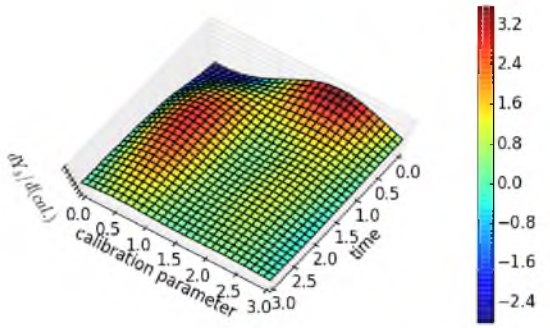

(a)

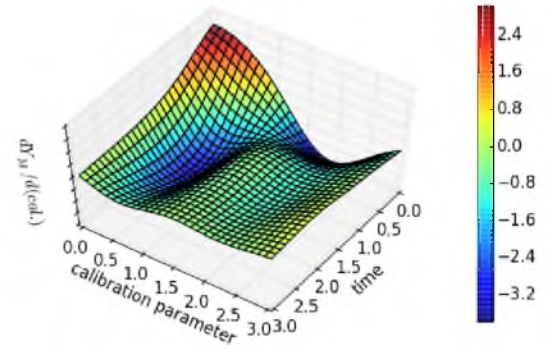

(b)

Figure 5.22: Three-dimensional view of discrepancy (a) and model GPs' (b) gradients with respect to the calibration parameter across time and calibration parameter space.

but by adding three more data-points (b) greater weight gets incorporated into the right mode of the distribution. The points removed were at times 0.754, 2.366, 3.010 and the points added were at times 0.266, 1.238, 2.527. Thus it appears that some portion of the nonidentifiability being observed can be attributed to the experiment being data poor.

Now to examine the effect that data quality has upon the current problem, the standard deviations in the Gaussian white-noise utilized to create the synthetic experimental datasets was altered, effectively creating more and less accurate data sampling. Figure 5.24 shows how decreasing the standard deviation to 0.15 (a) and increasing it to 0.45 (b) effects the calibration distribution. Reducing the variance in the experimental noise narrows the range of the calibration distribution, but creates a bimodal situation where both modes are of similar weight. Increasing the noise variance widens the range of the calibration distribution, but also significantly reduces the amplitude of the incorrect mode. The effect of changing the data quality on the models nonidentifiability does not appear conclusive

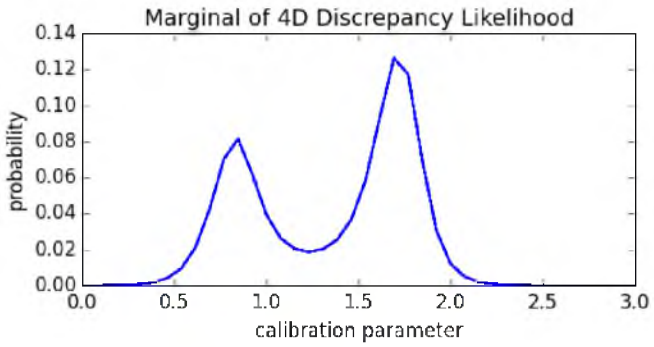

(a)

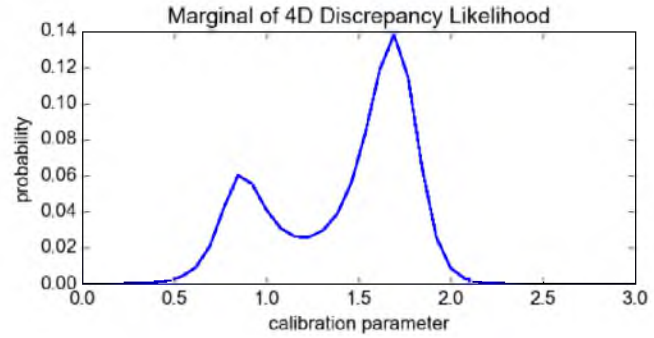

(b)

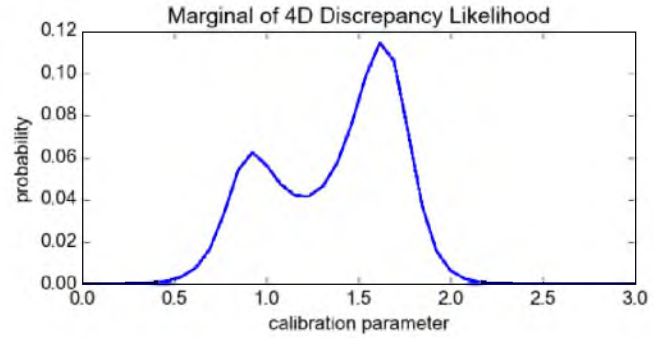

(c)

Figure 5.23: Three one-dimensional marginals of the discrepancy's four-dimensional likelihood distribution created in the same manner as Figure 5.13 except that that amount of experimental data available differs between plots. Compared to plot (c), which is the same as was shown previously, plot (a) used three less experimental data-points and the plot (b) used an additional three.

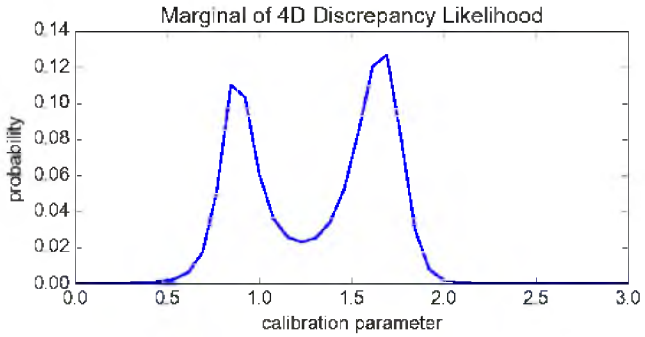

(a)

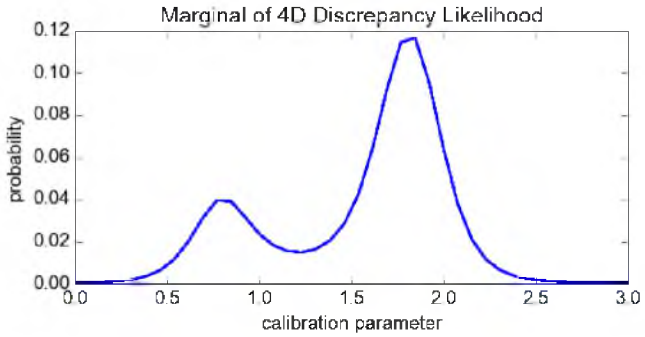

(b)

Figure 5.24: Two one-dimensional marginals of the discrepancy's four-dimensional likelihood distribution created in the same manner as Figure 5.13 except that that the standard deviation of the Gaussian white-noise used to generate the synthetic experimental differers between plots. Compared to Figure 5.23c, the plots were created with experimental data synthetically produced with a standard deviation of 0.15 and 0.45 in the Gaussian noise (as opposed to 0.3) for plot (a) and plot (b), respectively.

in this instance. It might be postulated that better calibration appears to occur with less precise experimental data because it allows the framework more leeway to balance the distribution of uncertainty between model-form and the experimental data.

5.5 Additional Approaches Towards Calibration

The Bayesian field of study is full of possibilities for tweaks and alterations of methods to accomplish new goals. The following three subsections describe a few ideas branching from Bayesian methods that were not found in a literature search during the summer of 2014. The three topics explored were Bayesian model comparison, multi-input Gaussian processes, and constrained Gaussian processes. These methods were considered as means of aiding calibration, thus potentially reducing the identification issue.

5.5.1 Bayesian Model Comparison

A Bayesian method of comparing model-forms and hypotheses is known as Bayesian model comparison. To compare two potential models (H_1 and H_2), a ratio of their probabilities described by Eq. (5.1) yields

$$\frac{P(H_1|D)}{P(H_2|D)} = \frac{P(D|H_1)}{P(D|H_2)} \frac{P(H_1)}{P(H_2)}. \tag{5.37}$$

This Bayesian comparison can be described as the posterior odds $\frac{P(H_1|D)}{P(H_2|D)}$ equaling the Bayes factor $\frac{P(D|H_1)}{P(D|H_2)}$ multiplied by the prior odds $\frac{P(H_1)}{P(H_2)}$. In cases where there are no prior preferences, or the prior odds are unity, Bayesian model comparison becomes a comparison of how well data can be explained by each hypothesis. It should be noted that Occam's razor, or favoring the least complex model that is reasonably able to explain the data, is explicitly contained within marginal likelihood of the GP. Referring back to Eq. (5.13), it can be seen that the first term on the right hand side of the equation $y^T K(X,X)^{-1} y$ encapsulates how well the data fit, while the second term $\log|K(X,X)|$ penalizes for model complexity [115].

Within Figure 5.25 the discrepancy's likelihood distributions, corresponding to calibration values uniformly sampled across the prior range of calibration values, are divided by the sum of all of the other likelihoods under consideration,

$$P(y|X, K(X,X), u_i) / \sum_i (P(y|X, K(X,X), u_i). \tag{5.38}$$

Within this method each Bayes factor calculated is effectively comparing H and $\bar{H}$, where $P(H) + P(\bar{H}) = 1$. From comparing the relative Bayes factors over the range of calibration values, it can be seen the values less than ≈ 0.7 and greater than ≈ 2.1 are unlikely. Although the peak around 1.8 clearly has the highest relative Bayes factor, other possible

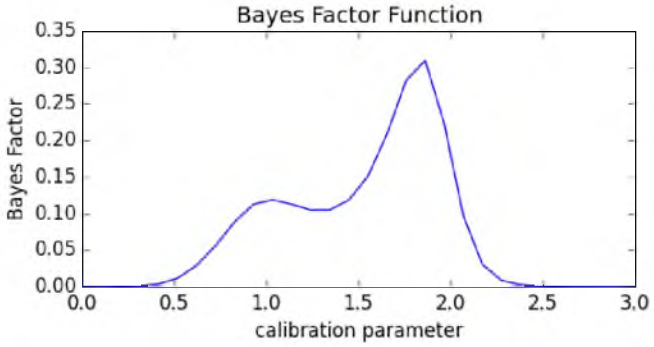

Figure 5.25: Displaying functionality of Bayes factor calculated as marginal likelihood over sum of other possible marginals.

values are within a 3:1 ratio of this maximum value. Within information theory and Bayesian equivalents, a ratio of greater than 3:1 is generally required to determine that a hypothesis is sufficiently supported to be more likely than another [72].

Bayes factors can be expressed in an evidence form, similar to those commonly seen in information theory, expressed as $10\log_{10}\left(\frac{P(D|H_1)}{P(D|H_2)}\right)$. This form is advocated by Jaynes (2003) [66], due to it being in terms of decibel units (db) that humans often find more intuitive. In decibels a 10 db difference corresponds to a factor of 10 difference [66]. Figure 5.26 shows the equivalent of Figure 5.25 in db units (a), as well as the equivalent for the H_2 discrepancy marginal likelihood (b) that will be further discussed in the following section. Although the same regions of the calibration space appear to be improbable, as was previously found by Figure 5.25, the degree of likelihood of each calibration, especially those of low probability, are easier to quantify.

5.5.2 Multi-Input Gaussian Processes

Inclusion of multiple data types is a means of aiding identification suggested by Arendt et al. (2012) [7]. To investigate this idea, H_2 data will be generated and incorporated into the analysis. First, H_2 models and experimental data are created in the same manner previously used for SiH_4. The incorrect model-form for H_2 will be

$$y_{\mathrm{H_2}}(t) = 2y_{0,\mathrm{SiH_4}}\big(1-\exp(-u\,t)\big), \tag{5.39}$$

while its corresponding 'true' model-form is

$$y_{\mathrm{H_2}}(t) = 2(y_{0,\mathrm{SiH_4}}-c)\big(1-\exp(-u\,t)\big). \tag{5.40}$$

From the true model, synthetic experimental data can be generated by taking samples including Gaussian white-noise with a standard deviation of 0.3 and three repetitions (Table 5.3), Eq. (5.40) + $\mathcal{N}(0, 0.3)$. The models with incorrect form, true form, and the experimental data sampled for both datasets are shown in Figure 5.27. Now that two datasets are available, a means of gaining more insight into the calibration term can be derived from probability theory.

Previously, when calibrating with a single data source the likelihood was only conditioned on that single dataset,

$$P\big(u, hp|y, X, K(X,X)\big) = P(y|X, K(X,X), u, hp)\; P(u, hp). \tag{5.41}$$

If multiple datasets are available, the increased amount of information contained within the likelihood should allow for the calculation of a more informed posterior. In order to

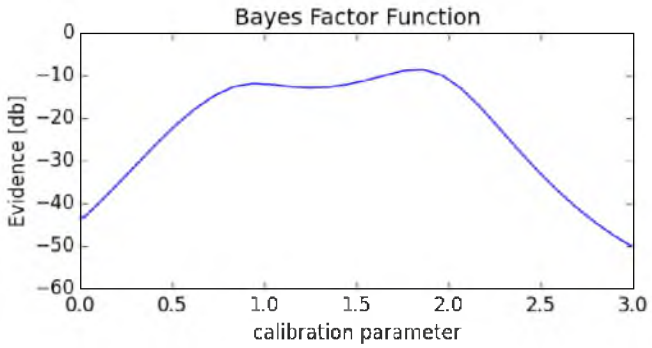

(a)

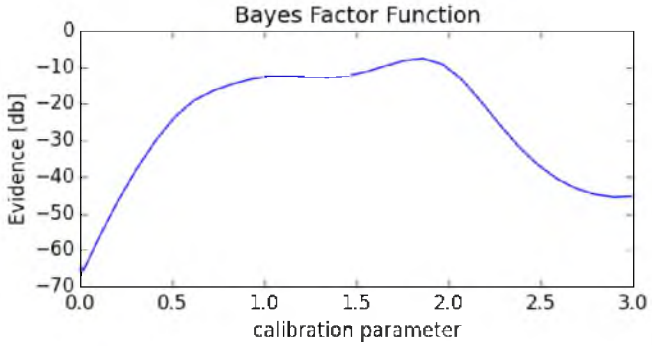

(b)

Figure 5.26: Functional behaviors of Bayes factors, equivalent to Figure 5.25, but now in terms of evidence as defined by Jaynes (2003) [66] for δ_{SiH_4} (a), equivalent to Figure 5.25, and δ_{H_2} (b).

Table 5.3: Experimental data times and corresponding three repetitions of H_2 concentrations.

	H_2 Concentration									
Time	0.110	0.432	0.754	1.077	1.399	1.721	2.043	2.366	2.688	3.010
y_1	0.616	3.306	5.197	6.556	6.155	6.656	6.602	6.786	6.921	7.010
y_2	1.411	3.840	5.156	5.498	6.099	6.909	6.961	6.610	6.974	7.242
y_3	1.047	3.585	4.803	6.196	6.370	6.866	6.857	6.474	6.868	6.885

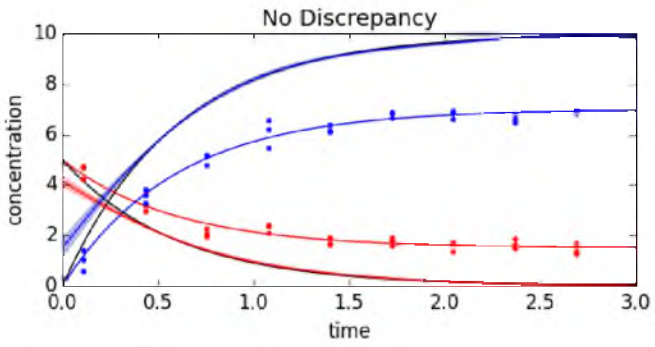

Figure 5.27: Incorrect models (black solid lines), GPs of incorrect models (shaded regions for 1 std. centered on colored line for mean), correct models (colored lines), and experimental data generated by sampling the correct models with added Gaussian noise (dots) for SiH_4 (red) and H_2 (blue).

create such a joint likelihood Eq. (5.42), the likelihood of the first dataset conditioned on the second dataset can be multiplied by the likelihood of the second dataset,

$$P\big(u, hp|y_1, y_2, X, K(X,X)\big) \propto P(y_1, y_2|X, K(X,X), u, hp)\ P(u, hp) \tag{5.42}$$

$$\propto P(y_1|y_2, X, K(X,X), u, hp)\ P(y_2|X, K(X,X), u, hp)\ P(u, hp). \tag{5.43}$$

This probabilistic expression can be captured within the scope of the current approach as follows (additional terms such as the inputs X, covariance kernels K and hyperparameters are assumed and noted as ...)

- First, GPs acting as surrogate models for both the incorrect model-forms can be created, $P(y_1^M|t, u, ...)$ and $P(y_2^M|t, u, ...)$.
- Then a GP modeling the discrepancy between the H_2 model GP and the corresponding experimental data can be created, $P(\delta_{y_2}|t, u, ...)$.
- Next, a GP modeling the discrepancy between SiH_4 and the corresponding experimental data can be created, but the concentration discrepancy data for H_2 that was utilized to create the H_2 discrepancy GP will be an additional input parameter, $P(\delta_{y_1}|\delta_{y_2}, t, u, ...)$.
- Once these two GPs have been created, their likelihoods can then be multiplied to effectively create the joint likelihood desired,
  $P(\delta_{y_1}, \delta_{y_2}|t, u, ...) = P(\delta_{y_1}|\delta_{y_2}, t, u, ...)\ P(\delta_{y_2}|t, u, ...)$.

- The joint likelihood can either be explored for the MLE calibration estimate or it can be multiplied by the hyperparameter prior to find the full posterior and MAP calibration estimate, $P(u|\delta_{y_1}, \delta_{y_2}, ...) = P(\delta_{y_1}, \delta_{y_2}|t, u, ...)\ P(u)$.

Within Figure 5.28 the likelihoods of the discrepancy GPs composed of only a single dataset are compared with the equivalent joint likelihood. Where both likelihoods based upon single datasets appear to have two major modes, the joint distribution has less of this characteristic. Although the joint likelihood's mode does not directly line up with the true value of the calibration parameter, it does greatly decrease the uncertainty within the distribution and provide greater confidence in its estimate. The joint likelihood can be calculated in two orders, exchanging which dataset is conditioned upon the other, and the two likelihoods resulting from the choice of ordering are shown in Figure 5.28c-5.28d. Probabilistically these joint distributions should be equivalent. Although the distributions shown are not completely equivalent, these two joint distributions are similar and differences are likely due to approximations made within the GPs. To create these four likelihoods, MCMC runs with 100 walkers were collected. The MCMC runs used 5,000 steps, with the first 2,000 discarded as the burn-in period. The hyperparameters of the surrogate model GPs were estimated by the MLE approach.

5.5.3 Constraining Gaussian Processes

There appear to be many potential uses of constraints within the KOH approach towards model-form uncertainty. Brynjarsdottir and O'Hagan (2014) [20] have demonstrated that constrained GPs allowed greater accuracy within interpolations, but did little to improve extrapolation. Da Veiga and Marrel (2012) [25] provide an overview of a variety of constraints that can be placed upon GPs including boundaries, monotonicity, and convex regions. Further details about integrating constraints into GPs are provided within Riihimaki and Vehtari (2010) [117].

When modeling chemical reactions, such as is the case for this exploration, physical constraints are easily identified. Such constraints include maximum and minimum concentrations, Eq. (5.44), monotonic trends (positive for products and negative for reactants), Eq. (5.45), and stoichiometric relations. Stoichiometry can be enforced through a comparison of the slopes of different chemical species traces, Eq. (5.46):

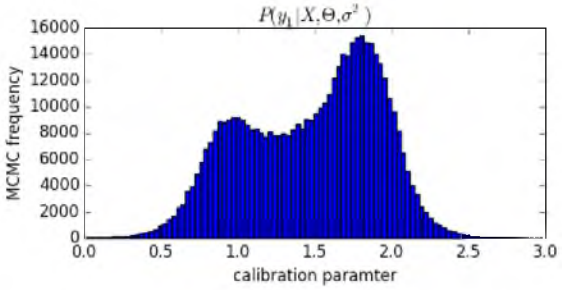

(a)

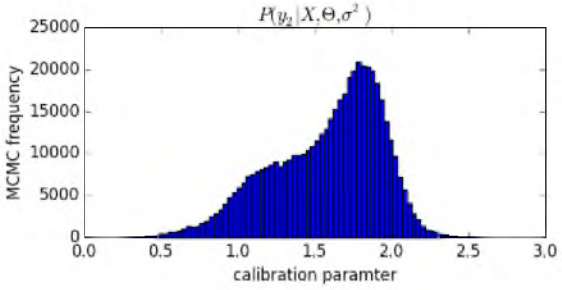

(b)

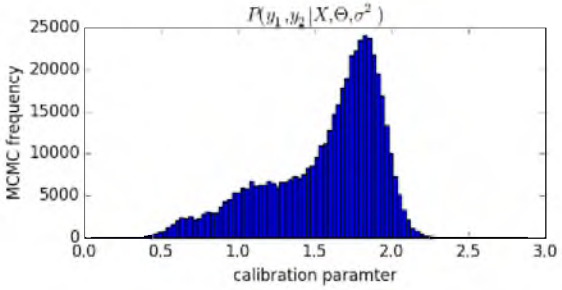

(c)

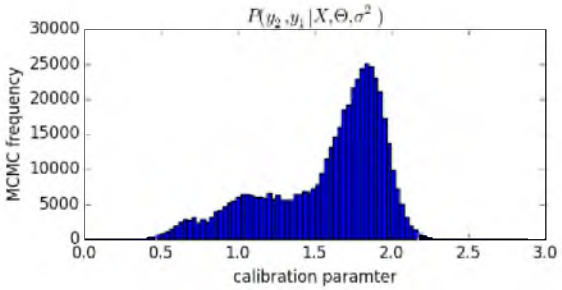

(d)

Figure 5.28: Comparison of likelihoods of single datasets' discrepancy GP, plots (a) and (b), with corresponding joint likelihood, plots (b) and (c). Likelihoods found by MCMC sampling of five-dimensional distribution for the likelihoods based upon single datasets and ten-dimensional for the joint likelihood. The joint was calculated in both orders, $P(y1|y2...)$ $P(y2|...)$ and $P(y2|y1...)$ $P(y1|...)$, as is reflected by the ordering in their titles.

$$\text{boundaries} \begin{cases} \text{if} & y_{SiH_4,i} > \text{bounds} \quad \text{for all } i & \rightarrow & \text{accept } \mathbf{y} \\ \text{else} & & \rightarrow & \text{reject } \mathbf{y} \end{cases} \tag{5.44}$$

$$\text{monotonicity} \begin{cases} \text{if} & \text{sign}(y'_{SiH_4,i}) = \text{expected sign} & \rightarrow & \text{accept } \mathbf{y} \\ \text{else} & & \rightarrow & \text{reject } \mathbf{y} \end{cases} \tag{5.45}$$

$$\text{stoichiometry} \begin{cases} \text{if} & -2\, y'_{SiH_4,i} > (1 - \text{tol})\, y'_{H_2,i} \quad \text{and} & & \\ & -2\, y'_{SiH_4,i} < (1 + \text{tol})\, y'_{H_2,i} & \rightarrow & \text{accept } \mathbf{y} \\ \text{else} & & \rightarrow & \text{reject } \mathbf{y}. \end{cases} \tag{5.46}$$

For the present pedagogical example, the absolute ratio of the slopes of SiH_4 and H_2 should approximately have a value of two. All three of these constraints were explored with varying degrees of success.

For the initial exploration of incorporating constraints, the GP acting as a surrogate model for the incorrect physics was manipulated. This entailed creating the GP (MAP estimated hyperparameters), collecting a ensemble of random samples from the GP, and then applying a pass/fail filter upon each individual sample. The constraints applied initially bounded the concentration between zero and five and placed a negative slope requirement between points on the sample trace corresponding to the fifth and sixth experimental data-points. This method effectively produced a truncated posterior distribution for the GP. An example of the effect these constraints had upon the posterior can be seen in Figure 5.29, where the posterior of the surrogate GP with calibration values of zero (a) and three (b) are shown. Filtering can be noted to shift the initial range of the posterior distribution down for the calibration value of zero, and shift the entire distribution up for the distribution with a calibration value of three. Both of the effects forced upon the GP distributions cause the GPs to be more physical. The mean of the truncated posterior was then compared to the mean of the experimental data in order to create the discrepancy GP. The effect of the constraints was then judged by its affect upon the marginal likelihood of the discrepancy GPs in terms of the calibration parameter, shown within Figure 5.30.

Although the constraints had a significant effect upon the surrogate GP, this did not translate into noticeable changes in the calibration marginal likelihood. It seems that this is likely due to the constraints primarily affecting regions where the calibration parameter took values already seen as low probability regions by the discrepancy likelihood, such as the two bounding calibration values used as examples. For more probable calibration values, minimal changes occurred due to the filtering, and thus they were not shown.

The stoichiometric filter was then implemented at two experimental data-points (second and forth) with an allowed tolerance of 15%. The effect that the stoichiometric constrained

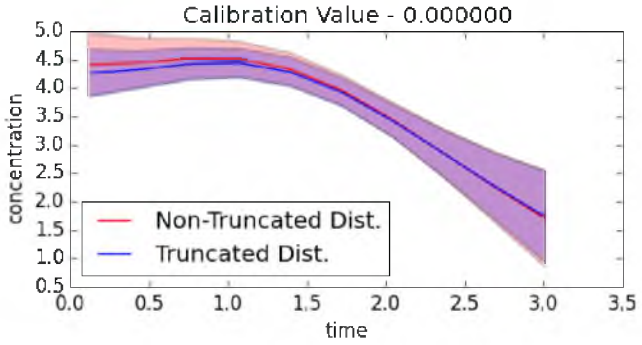

(a)

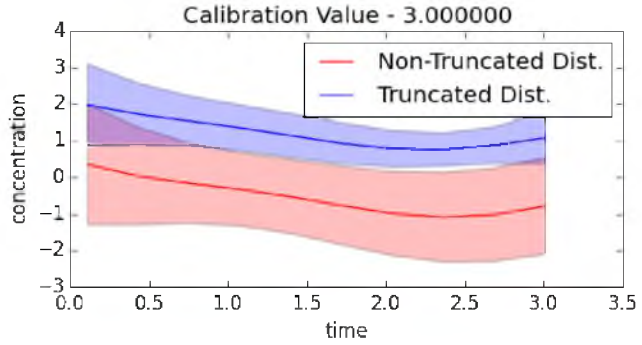

(b)

Figure 5.29: Posterior predictive distributions of surrogate GPs with and without constraints (truncated versus non) with calibration values of zero in plot (a) and three in plot (b). Fifteen hundred random samples were taken from the GPs and then a pass/fail filter bounding the samples between zero and five as well as a filter forcing a negative gradient between the fifth and sixth experimental data-points were applied.

surrogate GPs had upon the ultimate marginal posteriors for the discrepancy GP can be seen in Figure 5.30. Although there are not major changes in distribution appearance, the marginal corresponding to the constrained surrogate Figure 5.30b does have slight differences. The marginal with constraints has more weight on the second mode than the nonconstrained version, Figure 5.30a. Although the noted effect appears minor, it does offer potential as possible future avenue to be explored for improving calibration.

5.6 Forward Propagation

While the inverse problem has been focused upon thus far for the pedagogical example, ultimately the calibration distribution will be propagated forward through the model. The goal of such forward propagation of an uncertain parameter's distribution is to convert parameter uncertainty into predictive uncertainty for QoIs. Characterizing QoI uncertainty is a necessary step for any prediction. Figure 5.31a contains the calibration distribution

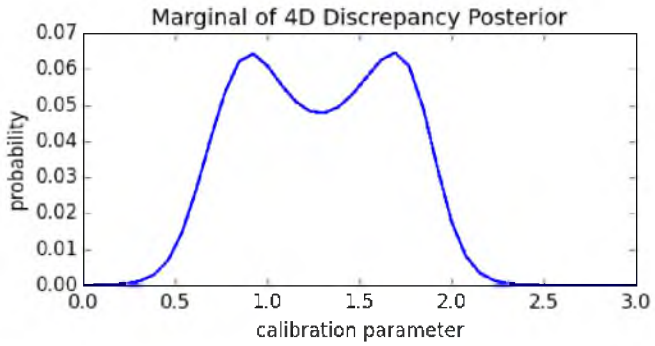

(a)

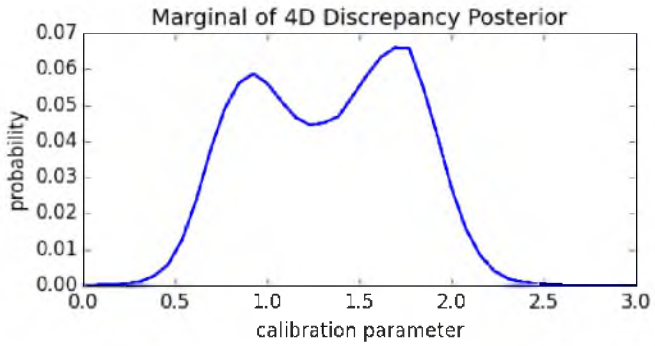

(b)

Figure 5.30: Comparison of marginals of four-dimensional posterior discrepancy GP, where surrogate GP was not constrained for plot (a), but was constrained for plot (b).

that was propagated forward through the surrogate and discrepancy models, as well as the QoI predictions produced with this forward propagation.

For this process, the surrogate model's hyperparameters were set using MLE and an MCMC exploration of the discrepancy GP was then utilized to create the calibration distribution. The MCMC consisted of 100 'walkers' taking 5,000 steps with a burn-in period of 2,000 steps. This calibration distribution was then propagated through the surrogate model and discrepancy by taking a screening the total MCMC sampling frequency of the calibration parameter, so that every 100th sample was utilized. This screening should remain representative of the whole distribution, while also removing any correlations between samples. For each calibration parameter value used in the discrepancy GP, the hyperparameter values were located with the MLE method. The resultant posterior predictive means are shown the left plot of Figure 5.31b and the mean and two standard deviations of those posterior predictive means are shown in the right plot. The true calibration parameter value lies within the region of high probability within this calibration distribution and thus will be used with high frequency when taking random draws.

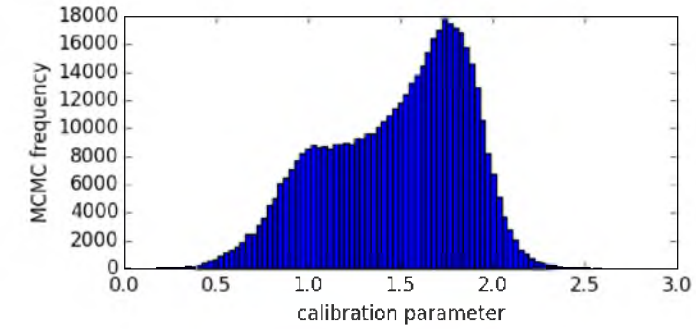

(a)

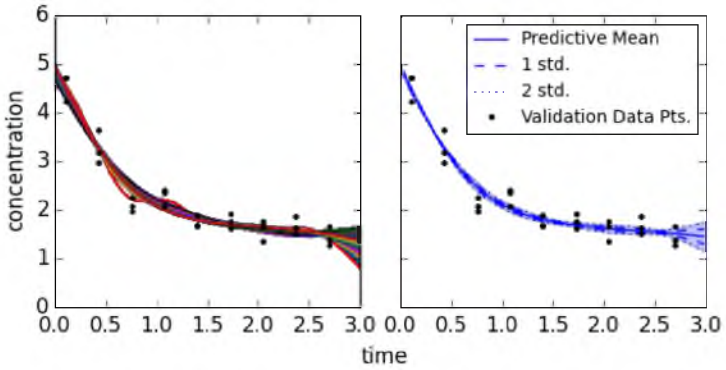

(b)

Figure 5.31: Forward propagation of calibration parameter distribution. The calibration distribution was created with an MCMC sampling frequency distribution of the discrepancy GP (a). Instances of posterior predictive mean traces created with instances based upon samples from calibration distribution are shown on the left in (b) and the mean and one/two standard deviations (solid and dashed/dotted lines, respectively, or the shaded region overall) of those instances are shown in the right plot of (b). Experimental data-points are included as black dots.

The posterior predictive distribution shown in Figure 5.31 demonstrates the effect of the calibration parameter uncertainty. Uncertainty in the GP itself is not included within this distribution because it was based upon the mean values of the aggregate of forward predictions. Two standard deviations of the posterior predictive distribution have relatively small variance compared to the equivalent in the experimental data. This small variance appears to be a result of the model-discrepancy negating any variance in the model outputs caused by the calibration uncertainty. Outside regions of experimental data such negation will no longer occur, as can be noted by the fast increase in variance as the time approaches the right temporal boundary. In order to better observe the effect of the calibration parameter uncertainty on the output QoI, the surrogate model GP's outputs can be directly observed in Figure 5.32.

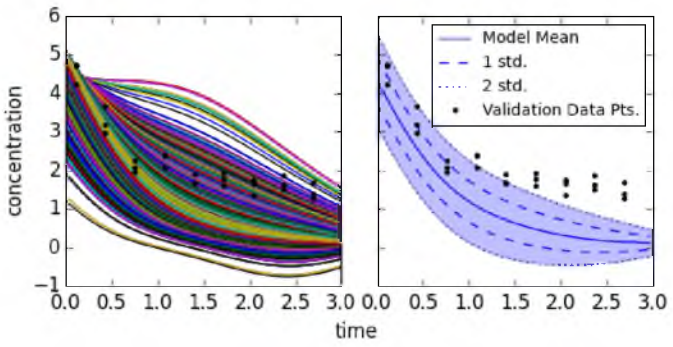

Figure 5.32: Forward propagation of calibration parameter distribution within the surrogate model. The calibration distribution was created with an MCMC sampling frequency of the discrepancy GP (Figure 5.31). Instances of mean value traces created with instances based upon samples from calibration distribution are shown on the left and the mean and one/two standard deviations (blue line and dashed/dotted lines, respectively, or overall shaded region) of those instances are shown in the right plot. Experimental data-points are included as black dots.

Forward propagation of the calibration parameter distribution through the surrogate model GP causes a much larger variance in QoI outputs than when combined with the discrepancy. Figure 5.32 shows the means of instances of GPs acting as surrogate models, where the calibration parameter was determined by sampling the distribution shown in Figure 5.31a. Just as in the previous figure for the posterior predictive, the left plot shows means of a sampling of the instances and the right plot shows the mean and standard deviations of those instance means. A single standard deviation of this distribution is of similar magnitude as the variance in the experimental data. In direct opposition to the effects noted for the variance in the posterior predictive, the variance shrinks in the right temporal boundary for uncertainty in the QoI. This is due to the fact that all the mathematical form of the model forces instances to converge to zero. Perhaps the most interesting information gained from this forward propagation can be found within its effects upon the model-discrepancy within Figure 5.33.

When using the KOH approach it seems that the model-discrepancy may be the best location for accessing model validation/consistency and the effects of uncertainty. The uncertainty in the discrepancy demonstrates the effects of the parameter uncertainty, while the mean discrepancy indicates bias error. By containing multiple forms of model error/uncertainty, comparison with experimental error should allow for thorough model validation and locating regions of model consistency. To better assess this, Figure 5.34 directly compares the discrepancy uncertainty with the experimental error.

Comparing the model discrepancy and its uncertainty due to the forward propagation

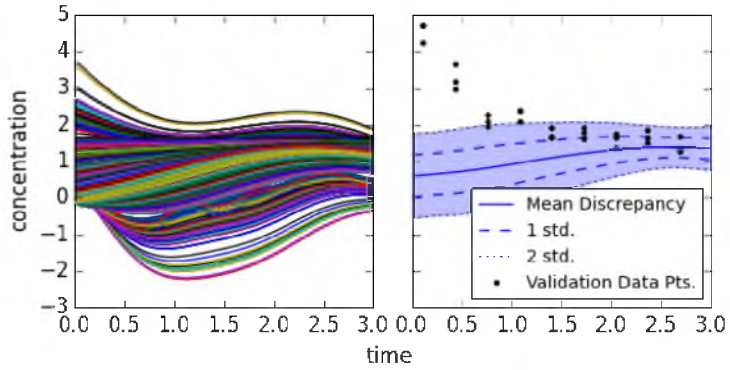

Figure 5.33: Forward propagation of calibration parameter distribution within model discrepancies. The calibration distribution was created with an MCMC sampling frequency of the discrepancy GP (Figure 5.31). Instances of mean value traces created with instances based upon samples from calibration distribution are shown in plot (a) and the mean and one/two standard deviations (blue line and dashed/dotted lines, respectively, or overall shaded region) of those instances are shown in plot (b). Experimental data-points are included as black dots.

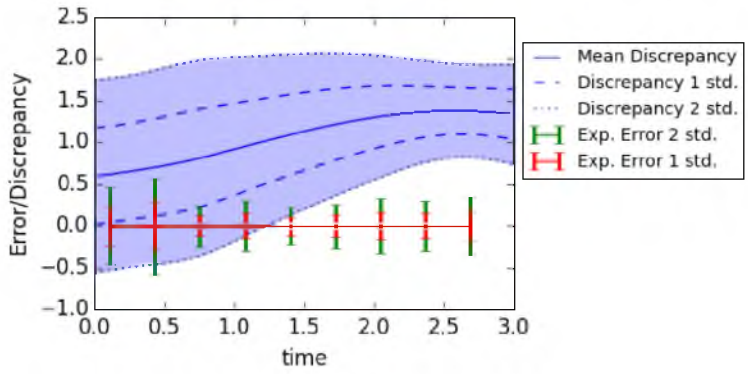

Figure 5.34: Comparison of uncertainty in model discrepancy due to forward propagation of calibration parameter uncertainty versus error in experimental data. Discrepancy mean and standard deviation, also shown in Figure 5.33, are shown as the blue region where the solid blue line indicates the mean value, the dashed line indicates one standard deviation, and likewise the dotted line indicates two standard deviations. The error bars indicate the standard deviations in the experimental data, where the red and green bars are for one and two standard deviations, respectively.

of the calibration distribution with the error in the experimental data provides a useful measure of consistency for the original model. In Figure 5.34 it can be seen that overlapping regions within a standard deviation of the discrepancy and experimental data exist. These overlapping regions are also near a discrepancy value of zero, meaning that predictions from the original model would be valid when using portions of the calibration distribution. The regions of overlap then shrink and move out to the second standard deviation of the discrepancy as time advances and the discrepancy shifts upwards. Near 1.5 time units there is no longer any overlap within two standard deviations of the discrepancy or experimental error. In the later time points, the uncertainty in the discrepancy is reduced, yet due to the previous upward shift there is still no potential consistent regions in the calibration distribution that could allow the model to overlap with the experimental data. While it is possible that calibration values of low probability, or those in the distributions tail, could allow for overlap, this would still indicate significant issues within the model or experimental data.

5.7 Conclusions

Throughout this chapter the Bayesian approach created by Kennedy and O'Hagan was explored and its boundaries tested. Only through application can a true understanding of the potential and limitations of a tool be understood. Clearly, Kennedy and O'Hagan's method of dealing with models containing uncertainty within the construction of their model-form is a valuable tool and can be applied to a wide array of problems. Interpolation with this methodology appears to be a good bet, while extrapolation is still something that must be approached with a critical mind. The identification problem, or simultaneous calibration and model-form uncertainty, is still an issue that must also be acknowledged whenever approaching this sort of problem. Model-form uncertainty is an area of research that will likely be a major focal point for the foreseeable future.

CHAPTER 6

COAL HEAT CAPACITY AND ENTHALPY SCALE-BRIDGING

6.1 Introduction

A scale-bridging model for coal particles' heat capacity and enthalpy are developed using Bayesian parameter estimation techniques. Scale-bridging involves the creation of a low-fidelity model to capture desired characteristics of a high-fidelity model at temporal and spacial scales appropriate for the application. The application for this model is large-scale computational fluid dynamics simulations of oxy-fired coal boilers. Scale-bridging of heat capacity and enthalpy is motivated by the desire to remove an iterative solve for particle temperature that currently exists within the code. Piecewise linear and piecewise quadric models are developed for this scale-bridging, followed by the gathering of experimental data and creation of probabilistic descriptions necessary to perform Bayesian calibration. Using a Markov chain Monte Carlo tool to perform the posterior distribution characterization, model-forms are compared, parameters reported, and demonstration provided.

6.2 Application

Within large-scale computational fluid dynamics (CFD) codes there are often opportunities to improve simulation speed. One such opportunity for reduced computational cost within the Carbon Capture Multidisciplinary Simulation Center's ARCHES code [139] is to improve the speed at which enthalpy and heat capacity of coal particles are calculated for coal boiler simulations. Particle enthalpy is a transported quantity within the CFD simulations, but particle temperature and heat capacity are not. Previously, these quantities were calculated using a model proposed by Merrick (1983) [95] that includes the evaluation of exponential functions and division operations. Each evaluation of particle temperature and heat capacity followed the following steps: propose a temperature for the coal particle based upon the previous time-step, evaluate the Merrick model for that temperature, compare the calculated enthalpy with the known (transported) particle enthalpy, and use an iterative

solve to converge. Once a particle temperature was found, Merrick's equation for heat capacity could be directly solved. While this method does locate the correct temperature, an iterative solve containing exponential function evaluations and division operations can become computationally expensive when implemented within large-scale CFD codes. Also, due to domain discretization that typically occurs within such codes, the control volume with the slowest iterative convergence effectively limits the speed of time-stepping for the entire simulation. To reduce this computational burden a scale-bridging model was developed.

6.3 Scale-Bridging Model Development

For this scale-bridging model design, the high-fidelity model upon which a low-fidelity model will be based will be the Merrick model [95]. Merrick (1983) presented a formulation for calculating sensible enthalpy h_{sens} and heat capacity C_p for coal particles. Mean atomic weight a is utilized throughout Merrick's equations and can be calculated by summing the ratio of elemental weight fractions composing the coal y_i and their respective atomic weights μ_i, ($\frac{1}{a} = \sum_i \frac{y_i}{\mu_i}$). This is useful because different coal types can have a wide variety of components and even for a single coal type the composition changes as the particle reacts. Two relations derived from the Einstein form of quantum-theory specific heat for solids are utilized within Merricks's equations,

$$g_0(z) = \frac{1}{\exp(z) - 1} \qquad g_1(z) = \frac{\exp(z)}{(\frac{\exp(z)-1}{z})^2}. \tag{6.1}$$

The sensible enthalpy h_{sens} and heat capacity C_p can then be calculated with

$$h_{\text{sens}} = \left(\frac{R}{a}\right)\left(380\ g_0\left(\tfrac{380}{T}\right) + 3600\ g_0\left(\tfrac{1800}{T}\right)\right) \quad [\text{J kg}^{-1}] \tag{6.2}$$

and

$$C_p = \left(\frac{R}{a}\right)\left(g_1\left(\tfrac{380}{T}\right) + 2\ g_1\left(\tfrac{1800}{T}\right)\right) \quad [\text{J kg}^{-1}\ \text{K}^{-1}], \tag{6.3}$$

where T is temperature [K] and R is the ideal gas constant [J/kmol K]. Three material subdivisions into which the solid particles are characterized are raw coal material, char, and ash. Eq. (6.2)-(6.3) can be used for raw coal and char by utilizing their respective mean atomic weights (i.e., accounting for all constituents of raw coal and pure carbon for char). For ash Merrick provides an additional relation

$$C_p = 754 + 0.586\ (T - T^0) \quad [\text{J kg}^{-1}\ \text{K}^{-1}] \tag{6.4}$$

from which

$$h_{\text{sens}} = 754\ (T - T^0) + \tfrac{1}{2} \cdot 0.586\ (T - T^0)^2 \quad [\text{J kg}^{-1}] \tag{6.5}$$

can be derived. Here $T^0 = 273.15$ is include as a reference temperature to convert degrees Kelvin to Celsius in order to maintain the model-form presented by Merrick.

The scale-bridging model will take a piecewise linear model-form for heat capacity and piecewise quadratic for enthalpy. These model-forms allow for a linear solve to be skipped within the application calculations, reducing computational costs. Previously a temperature was guessed, followed by an iterative solve for the enthalpy. Now using the piecewise enthalpy model, the temperature can be directly evaluated. As previously stated, enthalpy is a transported quantity within the CFD simulations. Given an enthalpy value the section of the piecewise model in which this value belongs can be located and then temperature solved for with a quadratic equation for that section of the piecewise model. Given the solved temperature, the piecewise heat capacity can be calculated. Formulation of the piecewise models occurs in the opposite order starting with the creation of the heat capacity model.

Two and three section piecewise linear model-forms for the heat capacity will be considered and compared:

$$C_p \cdot a = \begin{cases} C_0 + m_0 \cdot (T - T^0), & \text{if } T < T_1 \\ C_1 + m_1 \cdot (T - T^0), & \text{if } T < T_2 \\ C_2 + m_2 \cdot (T - T^0), & \text{else,} \end{cases} \tag{6.6}$$

where

$$C_1 = C_0 + (m_0 - m_1) \cdot (T_1 - T^0)$$

and

$$C_2 = C_1 + (m_1 - m_2) \cdot (T_2 - T^0).$$

Here C_i are intercepts, m_i are slopes, T_i are reference temperatures in degrees Kelvin. The two section model-form removes the last conditional statement in Eq. (6.6) and replaces the conditional requirement for the second section with an 'else.' There will be four and six random variables for the two and three section model variants, respectively. This piecewise linear model can be calibrated with data generated with Merrick's models [95]. The heat capacity multiplied by the mean atomic weight a is calculated, so that the equation can be scaled for raw coal, char, or any composition in-between. Utilizing the thermodynamic relation

$$\mathrm{d}h = C_p \, \mathrm{d}T, \tag{6.7}$$

the piecewise representation of heat capacity can be transformed into a quadratic form for enthalpy

$$\int_{h_i}^{h} \mathrm{d}h = \int_{T_i}^{T} C + m \cdot (T - T_i)\, \mathrm{d}T \tag{6.8}$$

$$h = h_i + C \cdot (T - T_i) + \tfrac{1}{2} m \cdot (T - T_i)^2. \tag{6.9}$$

This transformation allows enthalpy to be described by a piecewise quadratic model

$$h_{\text{sens}} \cdot a = \begin{cases} h_0 + C_0 \cdot (T - T^0) + \frac{1}{2} m_0 * (T - T^0)^2, & \text{if } T < T_1 \\ h_1 + C_1 \cdot (T - T^0) + \frac{1}{2} m_1 * (T - T^0)^2, & \text{if } T < T_2 \\ h_2 + C_2 \cdot (T - T^0) + \frac{1}{2} m_2 * (T - T^0)^2, & \text{else,} \end{cases} \tag{6.10}$$

where

$$h_1 = h_0 + (C_0 - C_1) \cdot (T_1 - T^0) + \tfrac{1}{2}(m_0 - m_1) \cdot (T_1 - T^0)^2$$

and

$$h_2 = h_1 + (C_1 - C_2) \cdot (T_2 - T^0) + \tfrac{1}{2}(m_1 - m_2) \cdot (T_2 - T^0)^2.$$

Now that piecewise models Eq. (6.6) and (6.10) are available, the path for solving particle temperatures is clearer. Given an enthalpy, h_i, its relative position among reference enthalpy values $h_{\text{sens}}(T_1)$ and $h_{\text{sens}}(T_2)$ is determined. Then the quadratic equation for that section of the piecewise model is solved, with the higher temperature solution always chosen due to the positive curvature of the enthalpy model. Then that temperature is compared to the reference temperatures T_1 and T_2 to determine the section of the piecewise heat capacity model to use and heat capacity is linearly solved.

6.4 Bayesian Calibration

Calibration of the scale-bridging model to the high-fidelity model will utilize Bayesian parameter estimation techniques. Prior to calibration, data from the Merrick model must be generated. Instead of simply using linearly spaced data across the temperature range of interest, prior knowledge of the particle temperatures typically found within the application boiler simulations will be incorporated. From boiler simulation data of the Boiler Simulation Facility provided by Dr. Benjamin Isaac of the University of Utah Carbon Capture Multi-Disciplinary Simulation Center, two particle distributions were estimated to approximate the typical particle temperature distributions found within the boiler. These distributions were a normal distribution $\mathcal{N}(1,650,130)$ and normal distribution $\mathcal{N}(1,100,130)$. A visual comparison of the two normal distributions with the original simulation data is provided

in Figure 6.1. The full synthetic dataset was created by combining 35 samples of the first normal distribution, 15 samples of the second normal distribution, and 50 temperature points linearly spaced from 300 to 3500 K. The additional uniform distribution was added to give the model full coverage of the temperature range of interest.

Bayesian parameter estimation is performed by solving Bayes law

$$P(X|y) = \frac{P(y|X)P(X)}{\int_{-\infty}^{\infty} P(y|X)\,\mathrm{d}X}, \tag{6.11}$$

where by multiplying prior knowledge of parameter values (prior distribution) $P(X)$ by inferred knowledge about the parameters from available data (likelihood distribution) $P(y|X)$, allows a probabilistic posterior estimate of the parameter value $P(X|y)$ to be obtained.

For this application it will be assumed that the uncertainty in the heat capacity values is Gaussian distributed, allowing the likelihood to also take a Gaussian form

$$P(y|X) = \frac{1}{(2\pi\sigma^2)^{n/2}} \exp\left(\frac{-\frac{1}{2}\sum_i^n (y_{M,i}(X) - y_{E,i})^2}{\sigma^2}\right), \tag{6.12}$$

where y_M are model outputs from the piecewise model, y_E are synthetic experimental data points, and σ is the standard deviation of the uncertainty in the experimental data. The standard deviation can be approximated for lower temperatures from data available in

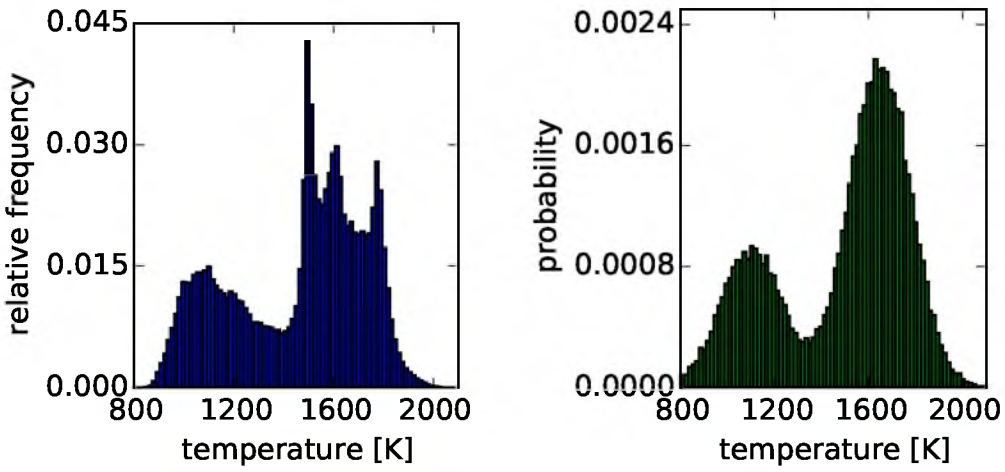

(a) (b)

Figure 6.1: Comparison of particle temperature distributions from a simulation with the distribution sampled for calibration data. Plot (a) shows the relative frequency of particle temperatures provided by Dr. Benjamin Isaac from a boiler simulation of the Boiler Simulation Facility. Plot (b) shows the distribution created to resemble plot (a), which was utilized for creating samples during parameter calibration.

Merrick (1983) [95] and then estimated as growing linearly from a threshold low temperature value or

$$\sigma = \begin{cases} 0.05 * 3R, & \text{if } T < 1000\,K \\ 0.05 * 3R + c_\sigma \cdot (T - 1000), & \text{else,} \end{cases} \tag{6.13}$$

where R [J / kmol K] is the ideal gas constant, 3 comes from a dimensionless scaling by Merrick, 0.05 is the estimate of error shown in Figure 2 of Merrick (1983) [95], and c_σ [J/ kmol K^2] is the estimated slope of uncertainty enlargement. Uniform prior distributions will be utilized for the uncertainty parameters, where prior knowledge will be used to define reasonable bounds: i.e., parameters determining temperatures at which the model shifts to a new section of the linear piecewise must be greater than zero but less than the next temperature bounds and slopes must be positive. Now that data, a likelihood function, and prior distributions are defined, a method of exploring the multidimensional posterior distribution can be utilized to provide distributions representing the uncertain parameters. The uncertain parameter distributions discovered for heat capacity can then be utilized within the enthalpy model.

To explore posterior distributions the Markov chain Monte Carlo (MCMC) Python tool emcee [38], previously discussed within Section 5.1.4, is used. After a few exploratory runs to locate good initial parameter values guesses, MCMC chains were initiated from the vicinity of those initial guesses by taking random 0.1% perturbations. For each MCMC run 100 Markov chains were evolved 3,000 steps and the first 1,000 were discarded as the 'burn-in' period to remove the influence of the initial guess.

6.5 Model Performance

While calibrating the model, two model-form considerations were explored: the number of sections comprising the piecewise models and the slope at which the uncertainty in the Merrick model increases with temperature. The model's performance with two and three sections were compared. Three sections appeared to be able to match the Merrick model well enough that additional sections would add complexity but minimally improve accuracy. Figure 6.2 shows marginal distributions from the MCMC results for a three section version of Eq. (6.6) with an uncertainty enlargement slope $c_\sigma = 1$. The distributions appear well defined, generally lacking bumps on the distributions that would be indicative of the 'burn-in' insufficiently removing the influence of the initial conditions. The distribution for m_2, plot (f), is clearly truncated at zero where its prior uniform distribution was bounded. Slopes less than zero would not be physical, so this truncation is acceptable. Most of the

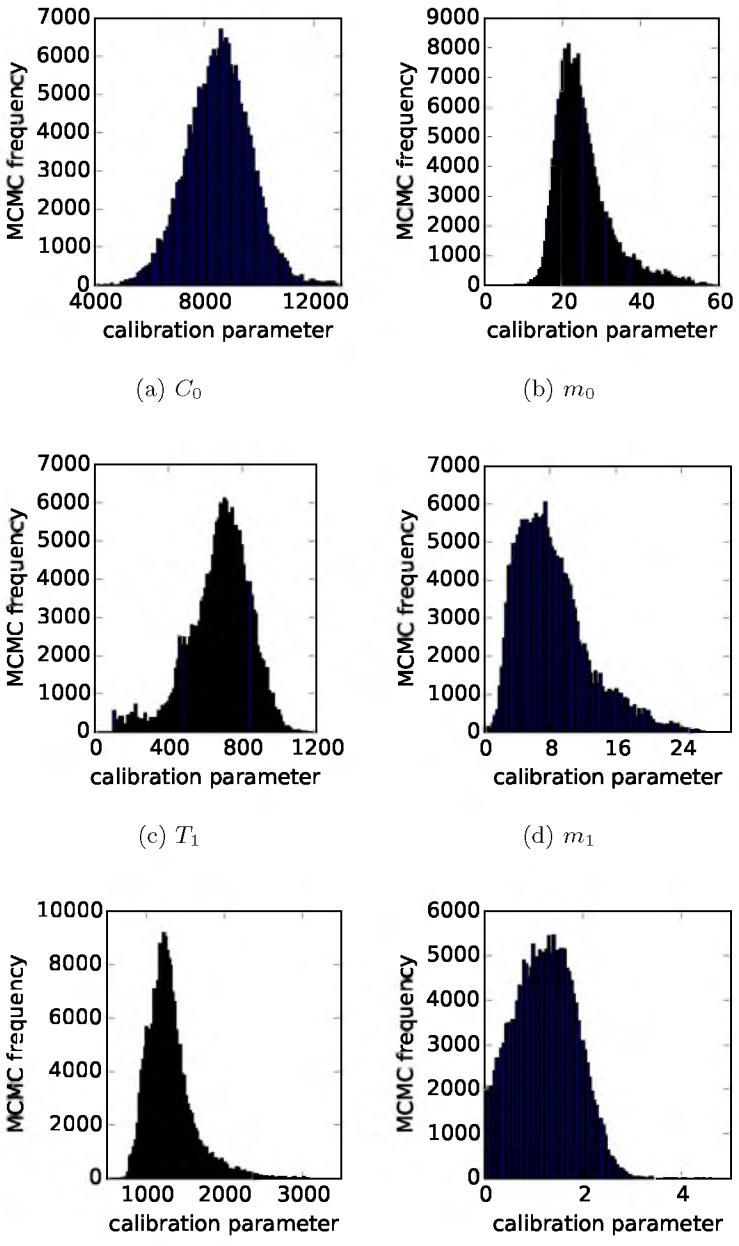

(a) C_0 (b) m_0

(c) T_1 (d) m_1

(e) T_2 (f) m_2

Figure 6.2: Illustrating the parameter estimation produced by MCMC exploration of posterior distribution for Eq. (6.6) where MCMC sampling frequency of parameter values reflects parameter posterior marginal distributions. The model variant explored within these plots was a three section piecewise linear C_p model where the experimental data's uncertainty increased with a slope of one.

marginal distributions appear to have non-Guassian form, indicating the multidimensional distribution has a complex, non-Gaussian form and that distribution means and modes will differ.

To visually compare the performance of the four model variants explored, Figure 6.3 shows the model's prediction for the temperature points at which experimental data were generated for calibration. Table 6.1 contains the marginal distribution means and full posterior mode for the uncertain parameters explored for each model variant. The effect of the number of segments composing the piecewise model is most clearly seen comparing plots (a) and (d) where the only difference between the model variants was changing the number of segments from three to two, respectively. As would be expected, the three section variant is a closer match to the Merrick model data, but if the uncertainty in the Merrick model is used to judge the model performance, the two section variant's only significant issue lies with the region of the model near 900 K. Near 900 K is where the two section variant switches sections and it is here that the model performs poorly in mimicking the Merrick model.

To compare the effect that the rate of increase in data uncertainty has upon the model's performance, plots (a, b, c) show model variants comprised of three piecewise-sections where the uncertainty was varied from 0.5 (b) to 1 (a) to 1.25 (c). It could be presumed that large slopes in uncertainty enlargement would cause the model's calibration to be weighed in favor of lower temperature data, yet the larger quantities of data points utilized at higher temperatures appears to have countered that intuitive result. Visual comparison of plot (a, b, c) yield little notable difference correlated to the uncertainty in the data. Delving into Table 6.1, modes from the multidimensional posterior distributions as well as the means of the marginalized distributions for each parameter do not show trends in the parameter values correlated to the changes in data uncertainty. Thus model-forms with $c_\sigma = 1$ will be utilized moving forward.

Also significant towards the evaluation of the scale-bridging models is the comparison of the enthalpy predictions with the Merrick predictions, as shown in Figure 6.4. The most probable parameter values for the model variants with $c_\sigma = 1$ were utilized to create the scale-bridging model results. To calculate the enthalpy, the h_{sens} calculated with Eq. (6.2) and (6.10) have a reference enthalpy, $h_{\text{ref}} = h_{\text{sens}}(25°\text{C})$, removed. For the scale-bridging model this is accomplished by solving for h_0 such that $h_{\text{sens}}(25°\text{C}) = 0$. For the scale-bridging model variant with three sections and $c_\sigma = 1$, $h_0 = -2.25E5$ J/kmol K when using the posterior mode parameter values. This reference enthalpy is removed

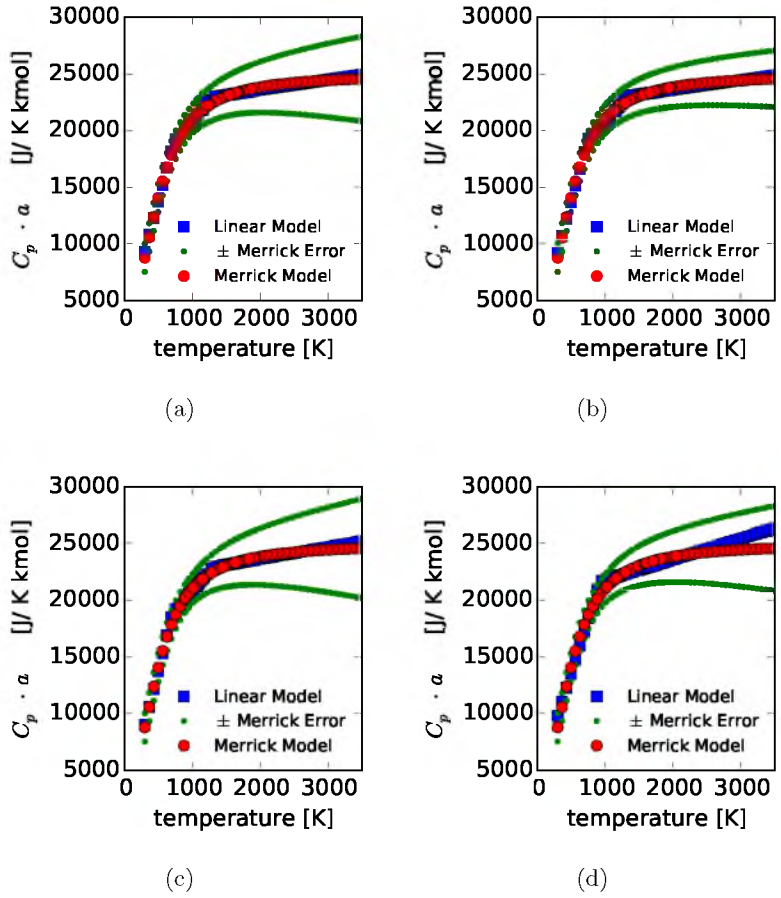

Figure 6.3: Comparing performance of variants of C_p piecewise linear model using mode parameter values found with MCMC searches where red circles indicate data samples calculated with Eq. (6.3), green dots are the red dots plus/minus one standard deviation of uncertainty c_σ and blue squares are produced by piecewise linear model Eq. (6.6). Plots (a, b, c) show variants of the piecewise model containing three sections while plot (d) shows a two section variant. The uncertainty standard deviation c_σ for plot (a) is 1, plot (b) is 0.5, plot (c) is 1.25 and plot (d) is 1.

Table 6.1: Mode from multidimensional posterior and mean values of parameter marginal distributions from the MCMC exploration of posterior distributions for parameter estimation. Order of parameters as they appear for each model variant left to right and top to bottom are C_0, m_0, T_1, m_1, T_2, and m_2.

Mode : Mean		
Linear 2 Section Piecewise Model with $c_\sigma = 1$		
9,290 : 9,310	18.9 : 19.2	927 : 924
1.83 : 1.84		
Linear 3 Section Piecewise Model with $c_\sigma = 0.5$		
8,600 : 8,480	22.7 : 24.5	738 : 706
6.26 : 7.46	1,360 : 1,370	0.841 : 1.02
Linear 3 Section Piecewise Model with $c_\sigma = 1$		
8,720 : 8,530	22.6 : 25.3	730 : 676
6.71 : 8.33	1,310 : 1,300	0.907 : 1.24
Linear 3 Section Piecewise Model with $c_\sigma = 1.25$		
8,340 : 8,550	24.2 : 25.1	710 : 679
6.69 : 8.39	1,310 : 1,330	1.04 : 1.26

so that the enthalpy of formation can be added within the application calculations. Both the enthalpy and heat capacity within Figure 6.4 are scaled by $a/3R$ following Merrick's example. From the figure it can be noted that the scale-bridging model's performance in matching the Merrick model's enthalpy predictions is not strongly effected by the number of segments utilized. The final step to using the derived models within the application calculations will be to scale the heat capacity and enthalpy by the mean atomic weights of the coal constituents. For the mineral matter Eq. (6.3) and (6.5) can still be utilized and are in the same mathematical form as the model's created for scale bridging, making vector implementations simple.

The last piece of analysis for this scale-bridging will be forward propagation of uncertainty associated with the calibrated parameters. Forward propagation of uncertainty allows the uncertainty in predictions to be estimated, providing quantification of predictivity. To achieve this, the parameter-sets accumulated by the MCMC algorithm's characterization of the posterior distribution can be used for evaluations of the piecewise models. After evaluating the piecewise models with the parameter-sets embodying the posterior distribution, mean predictions as well as other statistical characteristics of the model's predictions such as the standard deviation are available. Means and standard deviations for the enthalpy and heat capacity scale-bridging models are shown in Figure 6.5 for the two and three section variants.

Within the heat capacity plot it can be noted that the uncertainty propagated for both

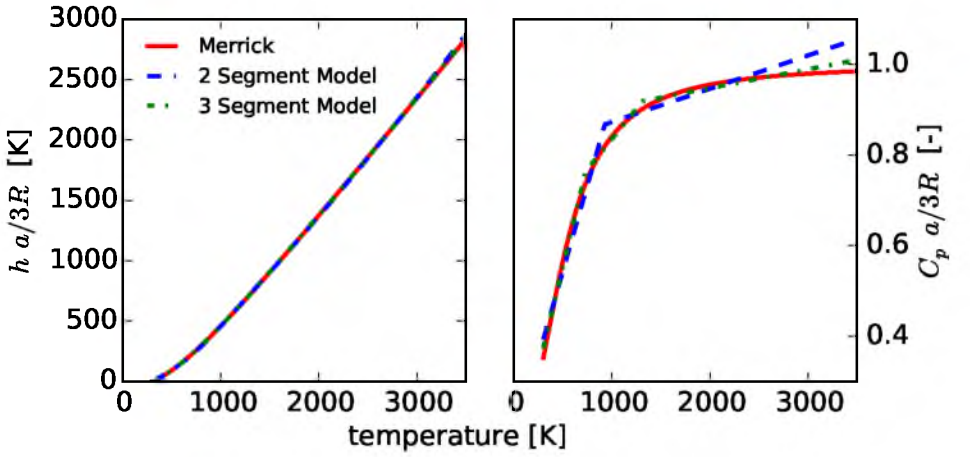

Figure 6.4: Comparing scale-bridging model performance for C_p and h_{sens} against the Merrick model. Red lines are from Merrick model, blue dashed lines are the two segment variant of the scale-bridging model, and green dash-dotted lines are the three section variant. The left plot shows h_{sens} scaled by $a/3R$ and the right plot shows C_p scaled by $a/3R$. For both scale-bridging model variants $c_\sigma = 1$. Mode parameter values were used for each model variant.

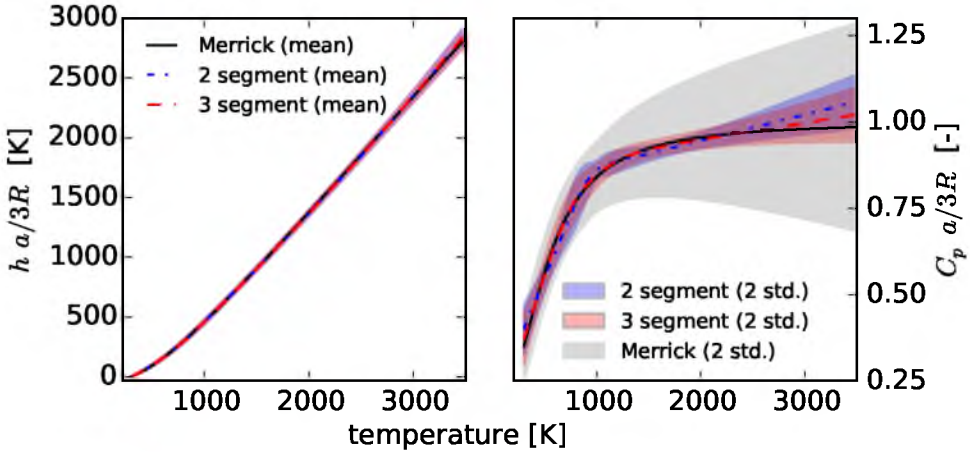

Figure 6.5: Forward propagation of parameter uncertainty from multidimensional posterior distribution. For both plots the blue dash-dotted line is the mean solution for the two section model, the blue shaded region signifies two standard deviations of the two section model's prediction, the red dashed line is the mean solution for the three section model, the red shaded region signifies two standard deviations of the three section model's prediction, the black line is the Merrick solution, and the grey shaded region shows two standard deviations of the uncertainty attributed to the Merrick model using $c_\sigma = 1$ within Eq. (6.13).

bridge-scaling model variants is far smaller than the prior uncertainty attributed to the Merrick model. Because the scale-bridging model's uncertainties are smaller than that attributed to Merrick, these model-forms will not be able to fully propagate uncertainty to applications, if that were desired. Means for both model variants are smoother than model evaluations based upon a single parameter-set, shown in Figure 6.4. The uncertainty for the two segment variant displays regions of higher uncertainty surrounding the piecewise segment connection, and less uncertainty at the connection point. The three segment variant's uncertainty generally appears smoother than the two segment variant and the Merrick model lies more within the three segment variant's uncertainty. The uncertainty associated with both model variants increases around the upper and lower temperature regions explored, as is expected due to less data available within those regions. The uncertainty for the enthalpy appears to be less effected by the scale-bridging model-form and increases with temperature. Although enthalpy's uncertainty appears minor, this uncertainty will effect the temperature that is utilized to calculate a heat capacity.

6.6 Conclusions

Piecewise linear and piecewise quadratic models for coal's heat capacity and enthalpy were created as scale-bridging models. These model-forms allow for a quadratic equation to be solved instead of an iterative solve involving equations with exponentials and division operations. This should allow increased computational efficiency for the application simulations. Using Bayesian parameter estimation with an MCMC tool, posterior parameter distributions were generated, allowing model variants to be compared. Upon comparing the performance of the model-forms, its appears that three section piecewise models perform well and the uncertainty in the synthetically created experimental data did not strongly influence the model's appearance. If the noted error in heat capacity for the two section model is acceptable, this simpler model could be utilized with little error propagated to the enthalpy calculation. This model can now be implemented and the portion of computational time spent on this piece of physics can be reduced with minimal effect on simulation outputs.

CHAPTER 7

COAL DEVOLATILIZATION SCALE-BRIDGING MODEL

7.1 Abstract

When performing large-scale, high-performance computations of multiphysics applications, it is common to limit the complexity of physics submodels comprising the simulation. For a hierarchical system of coal boiler simulations a reduced physics model is constructed to act as a scale-bridging method for a coal devolatilization model of greater physical rigor. This study is a variant of model-form uncertainty, with a more rigorous physics model, the chemical percolation devolatilization model, acting as a measure of the reduced physics model's performance for an application of interest. The application space is utilized to create a means of designing the reduced physics model by determining requirements and weighting desired characteristics. A single kinetic reaction equation with functional yield model and distributed activation energy is implemented to act as the reduced physics model. Consistency constraints are used to locate regions of the reduced physics model's parameter space that are consistent with the uncertainty identified in the rigorous model. Ultimately, the performance of the reduced physics model with consistent parameter-sets is validated against desired characteristics of the rigorous model and found to perform satisfactorily in capturing thermodynamic ultimate volatile yield trends and kinetic timescales for the desired application space. Framing the process of model-form selection within the context of validation and uncertainty quantification allows the credibility of the model to be established.

7.2 Introduction

Unlike more typical model-form uncertainty problems, where the goal is to the quantify a model's capacity to represent and predict experimental data, this application seeks to quantify a reduced physics model's (RPM) ability to stand in for a more complex, physics based model. The term model-form uncertainty will be used herein to refer to the difference

between a model prediction and reality. This idea was also referred to as model inadequacy within the seminal paper upon this topic by Kennedy and O'Hagan (2001) [74]. For the current application the RPM is calibrated to best fit desired characteristics of a more rigorous model for a given set of conditions. First and foremost the RPM must be able to produce results within the uncertainty of the more rigorous model's results. Ideally, the RPM will be able to capture the full range of the complex model's uncertainty, thereby allowing the full effect of the uncertainty in this set of physics be explored within the simulation applications. A RPM is based upon fewer parameters than the more rigorous model, and thus should allow easier propagation of the physics of interest's uncertainty into the ultimate applications. That being said, sensitivity to the RPM's parameters will be difficult to map back to parameters in the more complex model.

A process flow diagram of the approach taken herein for credible model development is shown in Figure 7.1. Initially the model will be tested for basic adequacy. In the terms of the flow diagram, adequacy is assessed through the comparison of simulation outputs with experimental data. Such comparisons are restricted to characteristics that can be quantified into comparable terms. The consistency evaluation determines if any subset of the parameter space allows the model to meet constraints specified for adequacy. Given that the model is shown adequate for the application, analysis of the model-form error, or discrepancy between model outputs and experimental data, will drive the model development towards validity. While this document focuses upon a specific application, the general concept of utilizing reduced physics models with quantified model-form uncertainty to bridge scales from more complex physics models is applicable to a wide array of engineering problems.

The process flow diagram also provides an overview of this document's organization. First the application of interest is introduced. Next, the experiments, or a more rigorous physics model within this implementation, is introduced and some of its characteristics such as input sensitivity and parameter uncertainty are analyzed. Following this will be a summary of the reduced model-form's development. The methodology used for consistency testing will then be defined as well as the application-based constraints upon which consistency is evaluated. Once consistent regions of the parameter space have been found, the validity of the current model-form is assessed and refinements to the model-form are proposed based upon the model discrepancy.

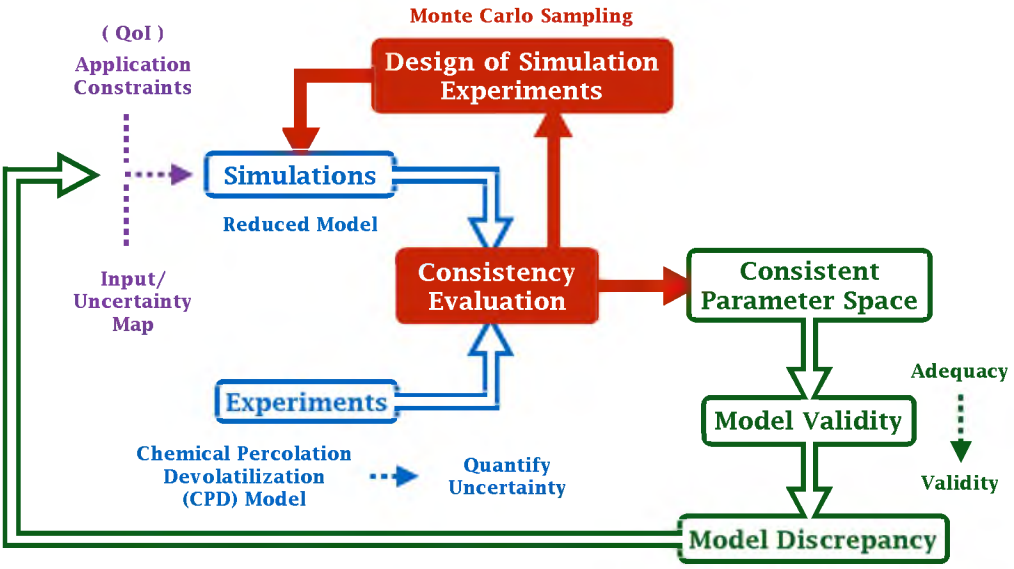

Figure 7.1: Process flow diagram of the method utilized to develop a scale-bridging coal devolatilization model.

7.3 Application Space

The University of Utah's Carbon Capture Multidisciplinary Simulation Center, a Predictive Science Academic Alliance Program (PSAAP) II Center, is primarily focused upon large-scale computational fluid dynamics (CFD) simulations of oxy-fired coal boilers. The project that the center has tasked itself with tackling is comprised of a hierarchy of high-performance computer-simulations spanning size-scales from the combustion of individual coal particles all the way up to a 350 MWe next-generation oxy-fire coal boiler. High-performance computer-simulations take significant computational resources, so simulation runs are limited. For the first year of the PSAAP II program, it was optimistically estimated that between eight and ten simulation runs per block of the hierarchy could be completed, thus limiting the validation and uncertainty quantification (VUQ) exploration of each of the blocks to two/three uncertain parameters within those confines. Each block of the hierarchy contains hundreds of uncertain parameters and tens of uncertain physics models. Focusing on two or three parameters from this list drastically limits what can be studied within a series of simulation runs. Uncertain parameters can be classified as scenario, model

or numerical. For the first year's simulations each block of the hierarchy was allowed to explore one model parameter and one/two scenario parameters. The scenario parameters were hierarchy-block specific including geometry specific attributes such as effective wall thickness or wall temperature profile, while the model parameter came from a piece of physics that ideally is shared among simulations.

For year one the piece of physics considered to have uncertainty to which the simulation outputs of interest will be most sensitive was coal devolatilization. Thus coal devolatilization physics will provide the model parameter for the first year's VUQ study. When studying coal devolatilization in this context, the application space cannot be ignored. Instantaneous function evaluations are necessary for the large-scale simulation, where devolatilization is calculated for each grid-point at each time-step. Devolatilization is the mechanism by which nonoxidized gases and tars move from the solid to gaseous phase for coal combustion or other thermal treatment, and thus have a significant influence upon the entire simulation. Complex models of coal devolatilization exist [131, 101, 37] and have been successful in describing experimental data, but are computationally too expensive to incorporate into large-scale CFD simulations. Simply put, a function with cheap evaluations that accurately captures the physical process was needed, or a scale-bridging model. Scale-bridging is a technique commonly found in simulation science [140, 91], where submodel complexity is limited by the simulation's resolution. Submodels are created to match the CFD's limitations while still capturing the desired physical characteristics.

Devolatilization is a chemical process that can be viewed from a kinetic or thermodynamic perspective. For the first year's simulations, capturing the ultimate volatile yield of the coal, or thermodynamic characteristic, is of first-order importance. The kinetics of how the coal reaches this ultimate yield will be considered of second-order importance. Kinetics are treated in this manner for the first iteration of simulations because devolatilization kinetics occur within a minimal region of the simulation domains and the amount of material transferred to the gas phase is judged to be more influential within the application than the rate at which this occurs.

7.4 The CPD Model

The 'true' physics of coal devolatilization is currently being represented by the chemical percolation devolatilization (CPD) model developed by Fletcher et al. (1992) [37]. CPD is an example of a rigorous devolatilization model that produces accurate results through relatively expensive evaluations. Within CPD, NMR spectroscopic data are utilized to char-

acterize the composition of a specified coal type. The CPD model includes a bridge-breaking reaction scheme, lattice statistics, percolation theory, and chemical phase-equilibrium calculations. CPD assumes a uniform temperature throughout the particle, or extremely small thermal Biot number, allowing the model outputs to be scaled by the mass of the desired particle size. In collaboration with Professor Fletcher of Brigham Young University, a MATLAB version of the CPD code [36] will be considered to produce 'true' physical data for this study. The CPD model predictions can be thought of as either 'truth' containing uncertainty or as a synthetic form of experimental data. Professor Fletcher provided uncertainty ranges for 13 uncertain model parameters within the CPD code, as shown in Table 7.1. These estimated ranges can be treated as uniform prior distributions based upon expert opinion. A secondary aspect of utilizing CPD for this study is that the credibility of the scale-bridging model will be leveraging CPD's credibility.

7.4.1 Temperature and Heating-Rate Effects

For developing models of coal devolatilization physics, two system conditions are typically considered as controlling the process: the rate at which coal is heated and the ultimate temperature that the coal reaches (hold-temperature). Due to the limitations of the ultimate application for the devolatilization model, a simplistic RPM was desired.

Table 7.1: CPD uncertain parameters and expert opinion uncertainty ranges solicited from Professor Fletcher of Brigham Young University.

Parameter	Nominal	Uncertainty	Max.	Min.
Coal Specific (Utah Sufco bituminous)				
p_0 [-]	0.483	0.03	0.513	0.453
c_0 [-]	0.0827	-	-	-
$\sigma + 1$ [-]	4.78	0.2	4.98	4.58
M_{clust} [kg/kmol]	457.8	20	477.8	437.8
m_δ [kg/kmol]	45.7	2	47.7	43.7
General CPD Model				
A_b [s^{-1}]	$2.6E+15$	5%	$2.73E+15$	$2.47E+15$
E_b [cal/mol]	55,400	5%	58,170	52,630
σ_b [cal/mol]	1,800	5%	1,890	1,710
ac [-]	0.9	0.05%	0.90045	0.89955
ec [-]	0	-	-	-
A_g [s^{-1}]	$3.0E+15$	5%	$3.15E+15$	$2.85E+15$
E_g [cal/mol]	69,000	5%	72,450	65,550
σ_g [cal/mol]	8,100	5%	8,505	7,695
A_{cr} [s^{-1}]	$3.0E+15$	5%	$3.15E+15$	$2.85E+15$
E_{cr} [cal/mol]	65,000	5%	68,250	61,750

Thus, the potential of eliminating one of these controlling system conditions from the ultimate model-form was considered. The initial investigation of the heating-rate and hold-temperature effects spanned a wider range of system conditions than was anticipated for the application, nevertheless this was utilized to gain a wider grasp of potential implications. Devolatilization is also sensitive to coal composition, but a Utah Sufco bituminous coal was used throughout the current analysis. Pressure can also effect this phenomena [37], but atmospheric pressure was assumed throughout this analysis because the boilers under consideration operate near atmospheric pressure. Nominal CPD parameter values were utilized through the hold-temperature and heating-rate analysis.

7.4.1.1 Hold-Temperature

To investigate the effect that the ultimate hold-temperature had upon coal's volatile yield due to devolatilization, CPD calculations were made over a range of hold-temperatures spanning 500 to 3,500 K, as is shown in Figure 7.2. These calculations all used a constant, linear heat-up rate of 10^9 K/s, which was assumed to effectively represent instantaneous heating to the hold-temperatures. Instantaneous heating of the coal was desired in order to isolate the hold-temperature's effect. The initial coal temperature was specified as 300 K. Once the hold-temperature was reached, the coal was held at that temperature for 10 seconds. This hold-time was assumed to be the infinite time-scale for the system of interest. For lower temperature and heating-rate systems, such an assumption would not be appropriate.

The vertical spikes in the volatile yield traces shown in Figure 7.2 are due to the coal reaching its specified hold-temperature and then continuing to produce volatile gases until reaching the effective equilibrium for that hold-temperature. The resulting equilibrium curve can be visualized in Figure 7.3, where data-points shown are the last points of the volatile yield traces in Figure 7.2.

It is evident that the ultimate volatile yield is strongly affected by the hold-temperature. The ultimate volatile yield curve represents an equilibrium curve for the application space of interest, but likely would change for applications in other domains such as underground in situ heating of coal where the heating-rates are much slower and the timescales are far longer. For the current application it appears that devolatilization initiates around 600 K and minimal changes occur above 1,600 K. While the asymptotic behavior for high temperatures can be debated, this behavior was not considered presently. Another use for this curve could be to act as a yield model. Volatile yield traces for temperatures above

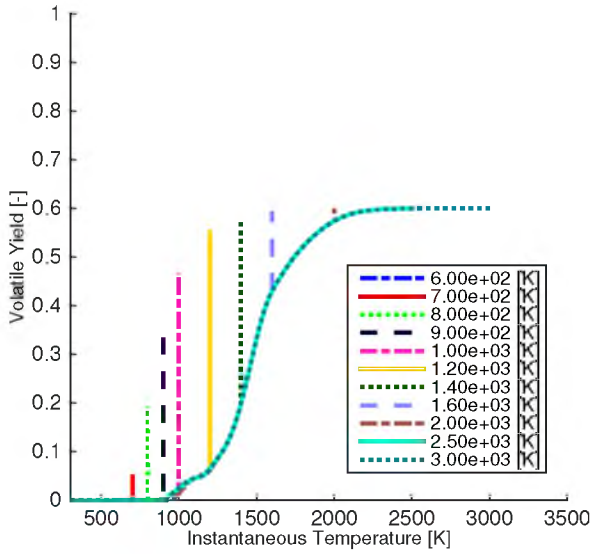

Figure 7.2: Volatile yield traces of Utah Sufco bituminous coal devolatilization for a range of hold-temperatures over their instantaneous temperature. To create the trace profiles the coal was linearly heated from 300 K to the specified hold-temperature at a rate of 10^9 K/s and then held at the hold-temperature for 10 seconds.

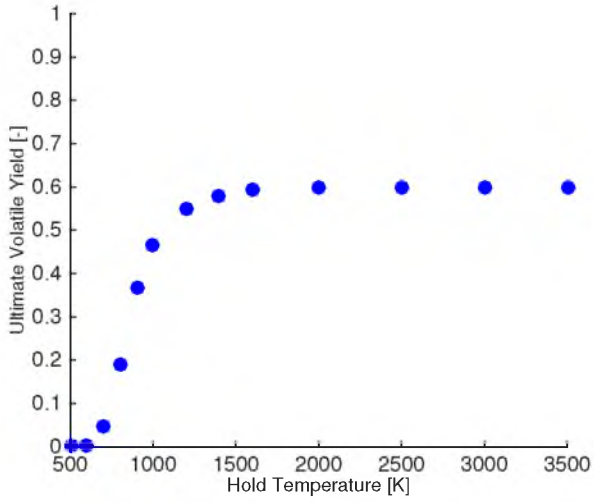

Figure 7.3: Ultimate volatile yield for Utah Sufco bituminous coal due to devolatilization over a range of hold-temperatures. Data points were extracted from traces shown in Figure 7.2 and were the product of linearly heating the coal from 300 K to the specified hold-temperature at a rate of 10^9 K/s and holding the coal at the hold-temperature for 10 seconds.

2,500 K were found to vary minimally, as can be noted in the overlap of traces for 2,500 and 3,000, and thus were not included in Figure 7.2.

Volatile yield is also a function of time, with coal exposed to lower temperatures losing volatiles over a longer period of time. This temporal functionality is shown in Figure 7.4. From the time traces it is clear that the chosen infinite time is in fact not infinite for coal at lower temperatures. For coal that reaches temperatures above 1,200 K, this infinite time assumption appears sufficient. Even with the now recognized deficiencies in this assumption, coal will not spend longer than 10 seconds in the boiler applications of interest, which allows this assumption to remain valid for this problem. All of the temporal traces above 2,000 K were found to strongly overlap and were left off Figure 7.4 for visual clarity.

7.4.1.2 Heating Rate

The effect of the heating-rate on the coal's volatile yield was examined in the same manner as the hold-temperature. Heating-rates ranging from 10^2 to 10^9 K/s were examined, as shown in Figure 7.5. To produce the traces shown, a hold-temperature of 1,600 K was utilized, which was reached after linearly heating the coal from 300 K. Again, a hold-time of 10 seconds was assumed to represent infinite time.

Heating-rates evidently strongly effect devolatilization kinetics. The traces for different heating-rates have different instantaneous temperatures for equivalent volatile yields, i.e., for

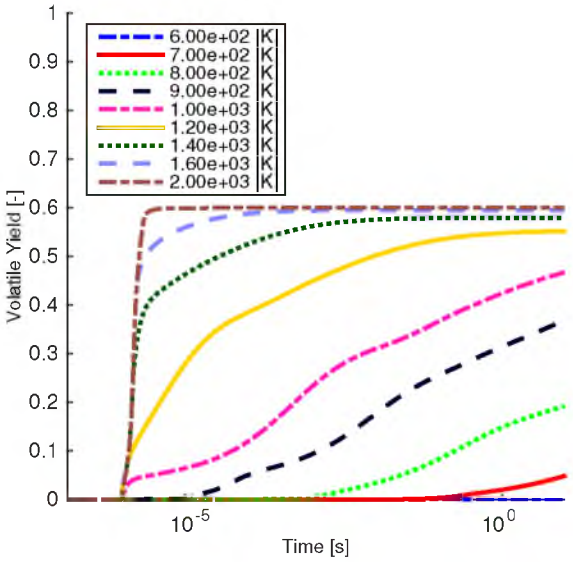

Figure 7.4: Volatile yield traces for Utah Sufco bituminous coal due to devolatilization for a range of hold-temperatures as they evolve over time. Coal was linearly heated at a rate of 10^9 K/s from 300 K to the hold-temperature and held at that temperature for 10 seconds.

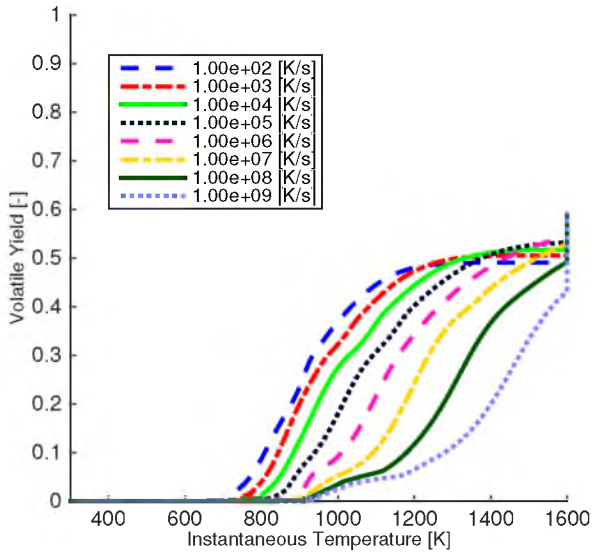

Figure 7.5: Volatile yield traces for Utah Sufco bituminous coal due to devolatilization for a range of heat-up rates over their instantaneous temperatures. These traces were calculated by CPD using a final hold-temperature of 1,600 K and were held at that hold-temperature for 10 seconds. For the range of heat-up rates shown, the coal was linearly heated from its initial temperature of 300 K to the hold-temperature.

coal heated at 10^9 K/s the instantaneous temperature for 20% volatile yield occurs around 1,400 K, while for coal heated at 10^2 K/s this occurs near 850 K. While the heating-rates demonstrate strong effects upon the kinetics, the effect upon the equilibrium curve shown in Figure 7.6 is less significant than was noted for the hold-temperature in Figure 7.3. To verify that the heating-rates are functioning as expected, Figure 7.7 was created to ensure the heating of the coal had approximately order of magnitude spacing.

7.4.2 Uncertainty Quantification

In order to thoroughly investigate the behaviors of CPD and eventually develop a RPM to approximate CPD, a design of experiments (DOE) was specified. The design was meant to cover regions of interest where CPD was believed to be accurate and/or which were important for the application boiler simulations. The grid-wise DOE covered five hold-temperatures and three heating-rates, as shown in Table 7.2. From what has been shown in Section 7.4.1, covering the range of hold-temperatures and heating-rates specified in the DOE demanded a dynamic model. Luckily, some of the uncertainty in CPD may reduce the difficulty of fitting these dynamics for the RPM.

The effect of uncertainty in the 13 uncertain CPD model parameters was characterized

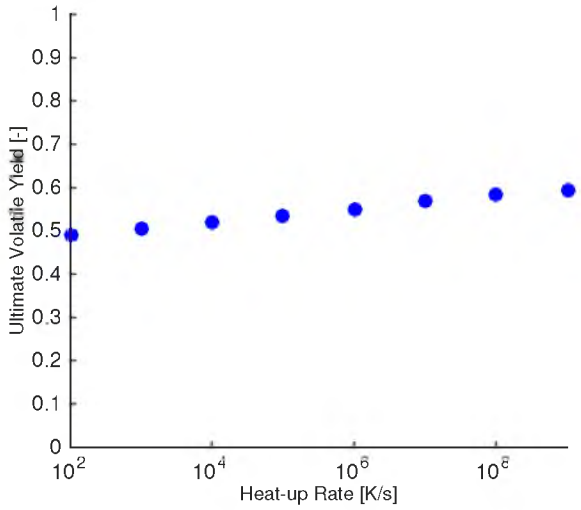

Figure 7.6: Ultimate volatile yield for Utah Sufco bituminous coal due to devolatilization over a range of heat-up rates. Data points were extracted from the data traces shown in Figure 7.5 and were produced by linearly heating the coal at the specified rates from 300 K to 1600 K and holding the coal at that temperature for 10 seconds.

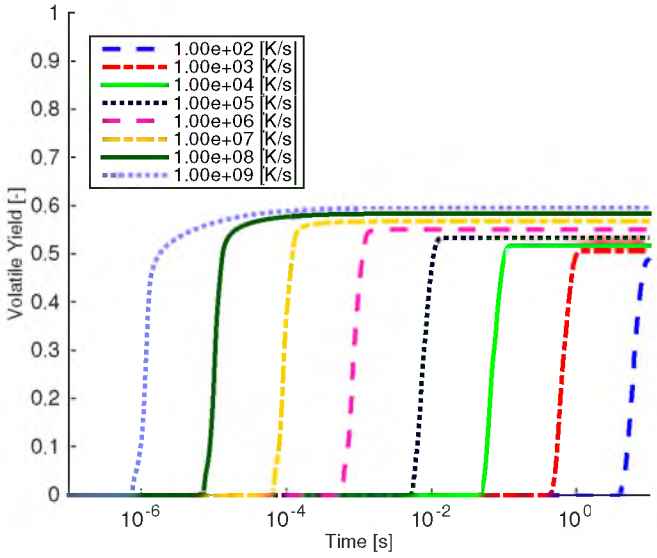

Figure 7.7: Volatile yield traces for Utah Sufco bituminous coal due to devolatilization for a range of heating-rates as they evolve over time. Traces created by linearly heating the coal at the specified rates from 300 K to 1600 K and holding the coal at that temperature for 10 seconds.

Table 7.2: Matrix of 15 design of experiments conditions.

Hold-Temperatures [K] : Heating Rates [K/s]		
700 : 1E4	700 : 1E5	700 : 1E6
1,000 : 1E4	1,000 : 1E5	1,000 : 1E6
1,300 : 1E4	1,300 : 1E5	1,300 : 1E6
1,600 : 1E4	1,600 : 1E5	1,600 : 1E6
2,400 : 1E4	2,400 : 1E5	2,400 : 1E6

with an uncertainty analysis completed for the 15 DOE points by taking 1,000 random samples from the uncertain parameter space, also known as a hypercube. A sample of the volatile yield trace's uncertainty produced by this random sampling is shown in Figure 7.8.

Clearly the parameter uncertainty causes significant effects within the CPD results. To better quantify the uncertainty in the CPD results, the data for the ultimate volatile yield was tabulated into statistical information within Table 7.3. The normalized standard deviation and normalized complete range provide quantitative measures of the direct effect of the model parameter uncertainties. Comparing those same statistics after marginalizing

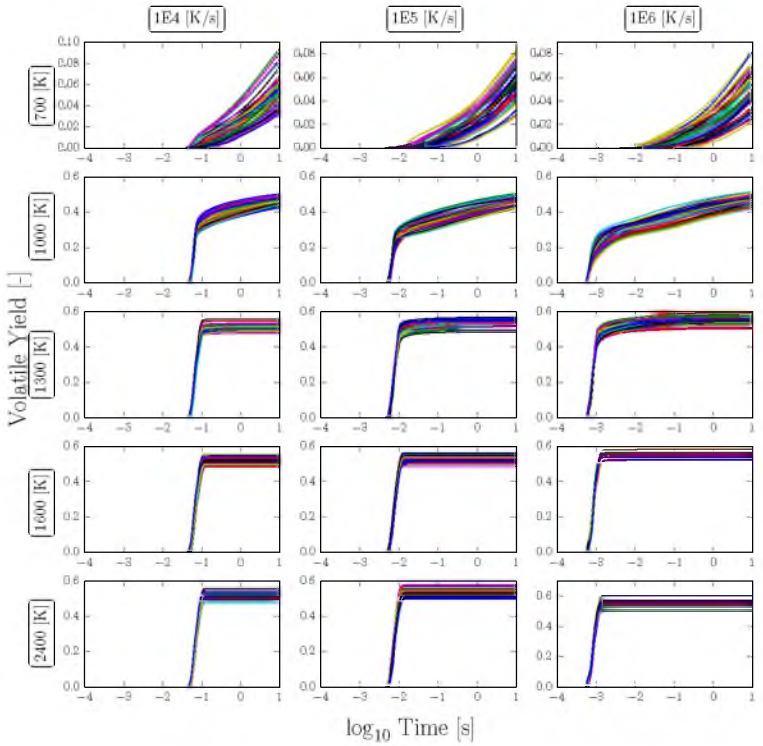

Figure 7.8: For each of the 15 DOE conditions, 50 Monte Carlo samples of the 13 uncertain model parameters used in the CPD calculations were utilized to create volatile yield traces demonstrating the effect of those uncertainties.

Table 7.3: Uncertainty statistics of 1,000 random samples of CPD's uncertain parameters arranged according to the DOE matrix in Table 7.2 and characterized in Figure 7.8. Values refer to ultimate volatile yield.

Full DOE								
Mean			Std.			Std / Mean		
0.0528	0.05348	0.05216	0.01540	0.01524	0.01576	0.2943	0.2945	0.3000
0.4605	0.4650	0.4652	0.02300	0.02331	0.02261	0.04974	0.05044	0.04875
0.5174	0.5325	0.5494	0.01908	0.01953	0.01881	0.03696	0.03670	0.03428
0.5174	0.5331	0.5504	0.01861	0.02010	0.01956	0.03593	0.03777	0.03557
0.5183	0.5329	0.5499	0.01908	0.01931	0.01888	0.03682	0.03623	0.03443
Min.			Max.			(Max. - Min.)/Mean		
0.02032	0.01922	0.01803	0.09839	0.09694	0.09776	1.493	1.5022	1.518
0.3991	0.4050	0.4058	0.5270	0.5328	0.5345	0.2771	0.2765	0.2776
0.4603	0.4754	0.4998	0.5696	0.5890	0.6078	0.2117	0.2135	0.1968
0.4686	0.4764	0.5023	0.5751	0.5852	0.6057	0.2057	0.2044	0.1880
0.4603	0.4843	0.4880	0.5696	0.5900	0.6068	0.2110	0.1982	0.2165
Marginalize Holding-Temperature								
[Max(Std/Mean) - Min(Std/Mean)] / Mean(Std/Mean)								
2.847	2.834	2.933						
[Max(Max-Min/Mean) - Min(Max-Min/Mean)] / Mean(Max-Min/Mean)								
2.683	2.722	2.774						
Marginalize Heating Rate								
[Max(Std/Mean) - Min(Std/Mean)] / Mean(Std/Mean)								
0.01932	0.03416	0.07459	0.1055	0.06037				
[Max(Max-Min/Mean) - Min(Max-Min/Mean)] / Mean(Max-Min/Mean)								
0.01678	0.003766	0.08035	0.08882	0.08746				

over the hold-temperature or heating-rate shows that the effect of hold-temperature is at least an order of magnitude more significant than that of the heating-rate. It should also be noted that the uncertainty in CPD parameters causes uncertainty in the volatile yields on the same order of magnitude as observed when varying the heating-rate.

7.4.3 Sensitivity Analysis

A baseline sensitivity analysis for the 13 uncertain model parameters in CPD was conducted exploring local and global sensitivities with a screening approach. Local sensitivities were estimated by changing uncertain parameters by 0.5% of their nominal value. This was done for one parameter at a time and the effects of adding and subtracting this perturbation were averaged, a one-at-a-time sensitivity measure variation [59]. Likewise, global sensitivity was assessed by moving one uncertain parameter at a time to the edge of its prior bounds (equivalent to taking an order of magnitude larger perturbation than the local sensitivity analysis). The local and global sensitivity of each uncertain parameter for all 15 DOE points was calculated in order to account for the hold-temperature and heating-rate effects

upon the sensitivities. The absolute amount that the ultimate volatile yield changes due to the altered parameter value is considered the sensitivity for this study. Sensitivities were normalized for each DOE point so that the largest is scaled to unity. This scaling allows simple identification of parameters having comparatively minor effect upon CPD outputs.

After analyzing the local and global sensitivities, five of the 13 uncertain-model parameters were deemed to contribute relatively minor amounts of uncertainty to the CPD calculations. Relatively minor is quantified as at least two orders of magnitude less sensitive for all DOE conditions than each individual condition's most sensitive parameter for both the local and global sensitivity. These five parameters were A_b, σ_b, A_g, σ_g, and A_{cr}. Thus, the majority of uncertainty in CPD calculations can be attributed to eight instead of 13 parameters. This allows less points to be tested when characterizing CPD's uncertainty, which ultimately was utilized to validate the RPM's performance. All five of the parameters that CPD is less sensitive to are general CPD model parameters, not coal specific, which makes this sensitivity finding applicable to studies for other coal types. It was noted that the local scaling of parameter *ac* was larger than the global scaling due to prior uncertainty bounds assigned, yet CPD was still found to be relatively sensitive to this parameter's uncertainty.

7.4.4 CPD Analysis Decisions

From the analysis of the CPD model useful knowledge has been gathered that can be used for future approximations. First, it appears that it is a better approximation to neglect the heat-rate effect than the hold-temperature when creating the RPM. Transporting the heating-rate history of each coal particle would have meant additional modifications within the CFD code, so the ability to make this assumption is useful. The other significant conclusion reached is that our primarily thermodynamic characteristic is only sensitive to eight of the 13 uncertain CPD parameters, so less sampling will be necessary to sufficiently sample the parameter space, or the 1,000 already collected are sufficient. The uncertainty space exploration will also be useful in defining the RPM's characteristic in the upcoming model development.

7.5 Reduced Physics Model-Form

For the application space of interest, the ultimate volatile yield is the first order effect of primary interest. When considering how to capture this effect, the approach of using a single first order reaction (SFOR) model, as presented by Biagini and Tognotti (2014) [14],

was considered,

$$\frac{\mathrm{d}V}{\mathrm{d}t} = A\exp(-E/T_P)(V_f - V). \tag{7.1}$$

Here A $[s^{-1}]$ is a preexponential factor, E $[K]$ is the activation temperature, T_P $[K]$ is the particle temperature, V_f $[-]$ is the ultimate volatile yield, and V is the volatile yield. Activation temperatures are equivalent to activation energies divided by the ideal gas constant. What differentiates this SFOR model from previous uses of single-reaction devolatilization models, such as that by Badzioch and Hawksley (1970) [9], is that the ultimate volatile yield is a function instead of a fixed value. Biagini and Tognotti proposed an exponential form with the ultimate volatile yield being a function of temperature

$$V_f = 1 - \exp\big(-DI \cdot \frac{T_P}{T_{st}}\big), \tag{7.2}$$

where DI stands for the composition specific, dimensionless devolatilization index and T_{st} is specified as the 'standard temperature' 1,223 K. This specific model-form did not closely resemble the thermodynamic yield curve produced by CPD, as shown in Figure 7.3, but further study of other forms appeared promising. Alternative functional forms can be explored through comparison with CPD results. Fitting

$$V_f = \frac{a}{2} \cdot \Big(1 - \tanh\big((b + c \cdot a) \cdot (590 - T_P)/T_P + (d + e \cdot a)\big)\Big) \tag{7.3}$$

to CPD data can be seen in Figure 7.9-7.10 where $b = 14.26$, $c = -10.57$, $d = 3.193$, and $e = -1.230$ for Utah Sufco bituminous coal. These values were found using a simplex minimization to fit the four curves shown in Figure 7.10 to the data extracted from CPD, as shown across both explored dimensions in Figure 7.9. Because the goal of this fitting was to match the CPD yield at low hold-temperatures, while allowing nonmatching, higher ultimate yields at high temperatures, only data-points with hold-temperatures less than or equal to 1000 K were used in the minimization. The value 590 K within Eq. (7.3) was physically meant to describe the temperature at which devolatilization begins. Although this volatilization temperature has currently been specified, it may be found to be coal dependent in the future.

From the three-dimensional view in Figure 7.9, it is evident that the yield model contained an error due to not capturing the effect of the heating-rate, especially at higher temperatures. Because the fitting included data from a range of heating-rates, the yield model was effectively fit to the center of the heating-rate range ($10^5 - 10^6$ K/s) and thus had the least error near those heating-rates. Overall, it can be noted that this yield model satisfactorily captures CPD's yield trend across a range of hold-temperatures.

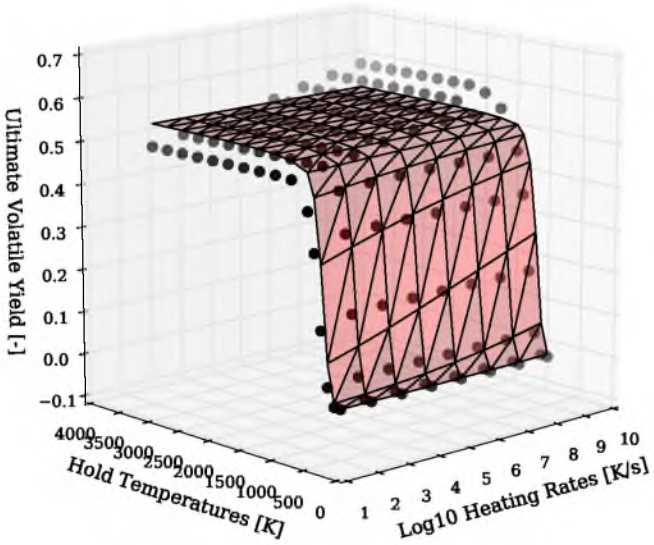

Figure 7.9: Fitting a yield model, Eq. (7.3), to CPD results across heating-rates and hold-temperatures. Dots signify CPD data-points and the surface is an interpolation of the fitted yield model.

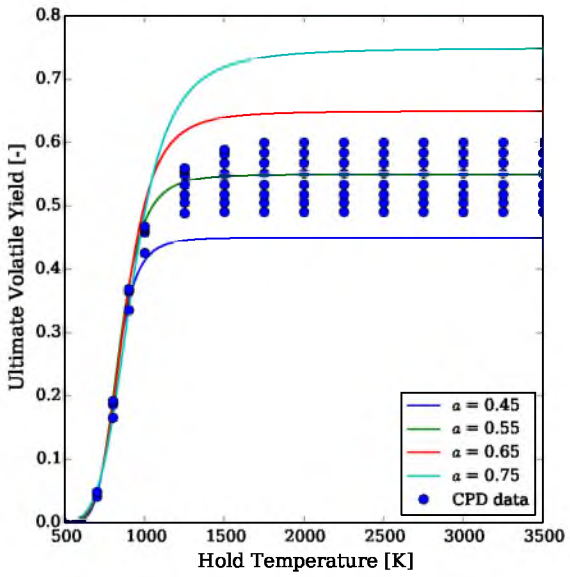

Figure 7.10: Demonstrating the effect of varying the high-temperature ultimate volatile yield parameter a on the yield model Eq. (7.3). Each line signifies a different a value and the dots are CPD data, also shown in Figure 7.9, where heating-rates and hold-temperatures were varied. CPD data contain vertical spread due to varied heating rate from $10^2 - 10^9$ K/s.

Figure 7.10 demonstrates an attribute of Eq. (7.3) for which this model was specifically designed. This model-form allows the ultimate volatile yield at high temperatures (relative to the application space of interest) be varied with the parameter a [−], while maintaining similar yields for lower hold-temperatures. This characteristic is allowed because of the sparsity of devolatilization data for higher temperatures. The optimal a value found through fitting with CPD data less than and equal to 1000 K was 0.54. Ultimately, this uncertain variable played a key role in exploring uncertainty in devolatilization for the application. Within Figure 7.10 the uncertainty in the parameter a can be noted to capture much of the same trend seen in the uncertainty within the CPD data due to the effect of the heating rate, with higher a values appearing to correspond with faster heating-rates, as well as the uncertainty in CPD outputs due to parameter uncertainty.

The simple SFOR model, Eq. (7.1), was unable to satisfactorily capture the desired physical characteristics of CPD for the specified DOE. In order to better reproduce the desired physical attributes, the concept of a distributed-activation energy model (DAEM) was incorporated into the reaction model. DAEM is based upon the idea of representing devolatilization as an infinite series of parallel reactions [5]. To model this concept it is assumed that there is a continuous distribution of activation temperatures and by evolving this distribution over time the effective activation temperature varies. The integral form of Eq. (7.1) with DAEM incorporated can be calculated as [129]

$$\frac{V_f - V}{V_f} = \int_0^\infty \exp\left[-\int_0^t k\,\mathrm{d}t\right] f(E)\,\mathrm{d}E, \tag{7.4}$$

where $k = A\exp(-E/T_P)$. A Gaussian distribution was assumed to describe the distribution of the activation energies or

$$f(E) = [\sigma_a(2\pi)^{1/2}]^{-1} \exp\left(\frac{-(E-E_0)^2}{2\sigma_a^2}\right). \tag{7.5}$$

One method of efficiently evaluating the DAEM model is to use the quadrature approximation to describe the distribution in terms of weights and abscissae [29]. Individual abscissae are evolved separately and then the final volatile yield is found by reapplying the weights and summing the then weighted abscissae. Alternatively, the DAEM can be approximated by the activation temperature distribution's inverse cumulative distribution function normalized by the full potential conversion or

$$Z = \max\big(-4.0,\ \min\big(\sqrt{2.0} * \mathrm{erfinv}(1.0 - 2.0 * (V_f - V)/a),\ 4.0\big)\big), \tag{7.6}$$

as illustrated in Figure 7.11 [37]. Note that the yield model's high-temperature ultimate yield a was used as the measure of conversion extent for this implementation. This distri-

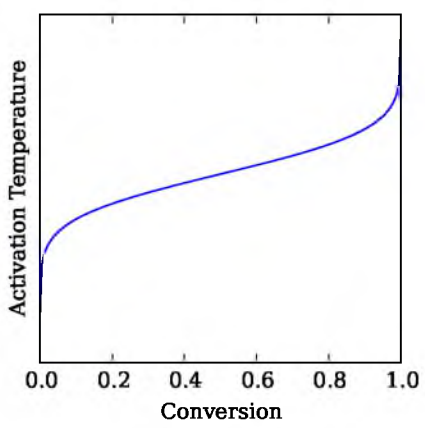

Figure 7.11: Activation temperature distribution's inverse cumulative distribution trace normalized to the amount of potential conversion.

bution has been truncated to avoid numerical issues with the tails of the distribution. With this method the effective activation temperature is initially low and as conversion progresses it increases. For devolatilization this causes reactions to initially accelerate quickly, but then decelerate while progressing towards completion. The inverse cumulative distribution representation of the DAEM was selected for implementation due to its simple computation.

Mathematically the ultimate form of the RPM was enacted as

$$\frac{\mathrm{d}V}{\mathrm{d}t} = \begin{cases} A \exp\left(\frac{-(E+\sigma_a Z)}{T_P}\right)(V_f - V), & \text{if } V_f - V \leq 0 \\ 0, & \text{else} \end{cases}, \tag{7.7}$$

where the activation temperature E $[-]$, the distribution's standard deviation σ_a $[K]$, the preexponential parameter A, and the yield model's high-temperature ultimate volatile yield a are free-parameters. The conditional is explicitly shown here not for the current calculations, but for the ultimate CFD application. Without this conditional statement in the formulation, reversed devolatilization could occur for particles moving from hot to cooler regions of the boiler. These four parameters can now be explored with a consistency analysis. Equation (7.7) is referred to hereafter as the Single Reaction with Yield Model (SRWY).

7.6 Consistency Evaluation

A first step in locating a set of parameters for the RPM, which can be utilized for the application, was to characterize a consistent space within the parameter's prior hypercube. The idea of a consistent space comes from the methodology described within Feeley et al. (2004) [34], where it was used to calibrate parameter values for the methane combustion

reaction-set, GRI-Mech 3.0. The basic premise of a consistent space can be described by the following equation:

$$(1-\gamma)l_i \leq M_i(\mathbf{x}) - d_i \leq (1-\gamma)u_i \qquad \text{for } i = 1, ..., N_{QoI}. \tag{7.8}$$

Here the model outputs, M, at specified parameter inputs, $\mathbf{x}$, are compared directly to experimental data, d. This comparison proves consistency for a particular quantity of interest (QoI) i if it is within the upper, u, and lower, l, error-bounds of the data. Model consistency is then achieved if consistency for all QoIs is found and the parameter ranges of the consistent sets overlap. This model consistency or inconsistency can be further characterized by the decimal fraction that the error bounds could be shrunk while maintaining consistency or the amount they could be expanded to reach consistency with the term γ. The parameter hypercube is specified as

$$\beta_p \geq x_p \geq \alpha_p \qquad \text{for } p = 1, ..., n, \tag{7.9}$$

where β and α designate the prior bounds for each parameter n. If model consistency is found, it will correspond to a subspace of the hypercube. If no consistent region is located, then the hypercube's prior bounds as well as the experimental error bounds could be reevaluated for possible expansion. A convenient method of approaching inconsistent systems is to look at unary and binary consistency or sensitivity to individual data-points for outliers.

For the current analysis the experimental data are produced by CPD. The uncertainty in the QoIs quantified by the uncertainty exploration (Section 7.4.2) will act as the error bounds. Random samples, such as those previously shown in Figure 7.8, were collected for use in defining QoI bounds. One thousand samples were taken for characterizing the uncertainty in CPD predictions. This quantity of sampling is deemed sufficient considering CPD was found to be sensitive to only eight of the 13 uncertain parameters (Section 7.4.3).

Defining QoIs is perhaps the most subjective component of the consistency analysis. The simplest method of selecting a QoI is to use the ultimate quantity in which a prediction is desired. An alternative selection procedure is to use features that relate to desired attributes deemed physical and/or necessary for accurate predictions of a physical phenomena that has no means of direct comparison. The two QoIs chosen to define consistency for the SRWY model are the ultimate volatile yield and the time to get to half the ultimate volatile yield. The ultimate volatile yield is the quantity directly desired from the SRWY model and represents capturing thermodynamic trends of coal devolatilization relevant to the application. Although kinetics were set to secondary importance when the heating-rate

effects were deemed less significant in Section 7.4.4, roughly capturing devolatilization kinetics is still desired. Thus, the second time-based QoI was selected to enable the SRWY model to roughly estimate the kinetic behaviors of CPD. The QoIs are visually depicted within Figure 7.12.

Using two QoIs for the 15 DOE points results in a total of 30 QoIs, as shown in Table 7.4, that must be simultaneously satisfied for consistency. Initially, all QoI values were strictly based upon uncertainty quantified in the uncertainty sampling of CPD, except for apparent outliers which were discarded. This was altered to accommodate greater perceived uncertainty that CPD does not take into account. There are high-temperature ultimate yield data reported to be higher than what CPD predicts [99]. An additional high-temperature data-point that could be considered is the sublimation point of graphite, which appears to have many caveats but is roughly estimated to be approximately 3950 K [2]. This potential discrepancy with CPD is believed to be due to such data not being taken into account during its formulation. In order to allow the SRWY model to reflect such high-temperature data, the ultimate volatile yield bounds for the DOE points with hold-temperatures 1,600 and 2,400 K were enlarged to reflect the perceived potential span. Through the course of exploratory consistency tests, it was deemed that the upper bounds of the ultimate volatile yield QoIs for DOE points at 1,300 K were limiting the higher temperature DOE point's ability to reach higher ultimate volatile yields with the current model-form. Thus, these bounds were also raised for all three heating-rates. The final

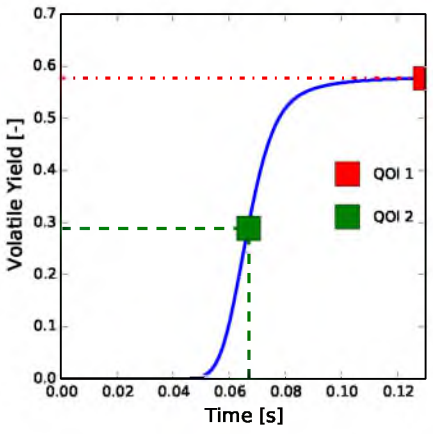

Figure 7.12: Illustration of QoI definitions utilized within consistency analysis. Blue line is a devolatilization volatile yield trace over time. The red and green boxes demonstrates the position of first and second QoI, which are the ultimate volatile yield and time to half the ultimate volatile yield, respectively. Dashed and dashed-dotted lines illustrate where the boxes fall on the volatile yield and time-scales.

Table 7.4: QoI ranges for 15 DOE points. QoIs are the ultimate volatile yield and the time to get to half the ultimate volatile yield.

QoIs			
	DOE 1	DOE 2	DOE 3
$\frac{1}{2}$ Ult. Vol. Yield Time [s]	8.099e-01 – 2.857e+00	7.544e-01 – 2.896e+00	7.627e-01 – 2.897e+00
Ult. Vol. Yield [-]	2.032e-02 – 9.839e-02	1.922e-02 – 9.694e-02	1.803e-02 – 9.776e-02
	DOE 4	DOE 5	DOE 6
$\frac{1}{2}$ Ult. Vol. Yield Time [s]	6.312e-02 – 7.019e-02	6.933e-03 – 9.686e-03	9.058e-04 – 3.728e-03
Ult. Vol. Yield [-]	3.991e-01 – 5.270e-01	4.050e-01 – 5.328e-01	4.058e-01 – 5.345e-01
	DOE 7	DOE 8	DOE 9
$\frac{1}{2}$ Ult. Vol. Yield Time [s]	6.402e-02 – 7.381e-02	7.035e-03 – 8.073e-03	7.867e-04 – 1.155e-03
Ult. Vol. Yield [-]	4.603e-01 – 6.500e-01	4.754e-01 – 6.500e-01	4.998e-01 – 6.500e-01
	DOE 10	DOE 11	DOE 12
$\frac{1}{2}$ Ult. Vol. Yield Time [s]	6.456e-02 – 7.415e-02	7.120e-03 – 8.068e-03	7.881e-04 – 1.153e-03
Ult. Vol. Yield [-]	4.500e-01 – 7.000e-01	4.500e-01 – 7.000e-01	4.500e-01 – 7.000e-01
	DOE 13	DOE 14	DOE 15
$\frac{1}{2}$ Ult. Vol. Yield Time [s]	6.443e-02 – 7.392e-02	7.076e-03 – 8.075e-03	7.848e-04 – 1.152e-03
Ult. Vol. Yield [-]	4.500e-01 – 7.600e-01	4.500e-01 – 7.600e-01	4.500e-01 – 7.600e-01

alteration to the QoIs was to increase the temporal QoI uppers bounds by multiplying them by 1.3 for the DOE points with 1E6 K/s heating-rate and hold-temperatures of 1,300, 1,600, and 2,400 K. These bounds were extended to allow greater amounts of consistency across all DOE points. With high temperatures and fast heating-rates, the kinetic timescales of those three DOE points should have minimal effect upon the application simulations.

Once the QoI definitions were established and values set, random samples of the free parameters in the RPM were tested for simultaneous consistency across all QoIs and γ values collected for consistent samples. Where typically γ values indicate the ability to shrink data error-bounds, for this model-form uncertainty application a more useful interpretation of γ is that it is indicative of how true the RPM fits the center of the QoI's uncertainty. In order to capture the full spectrum of CPD uncertainty, a range of γ values would be required.

Even with these relatively simple QoIs there is ambiguity about how they are defined and enacted. For each SRWY trace the ultimate volatile yield and time it took the SRWY trace to get to half of that ultimate volatile yield are compared with the CPD ranges of uncertainty for those same quantities. An alternative comparison that could be explored would be to check if the SRWY model's trace passed through the volatile yield range equal to half the uncertain CPD ultimate volatile yield range within CPD's time frame for reaching half its ultimate volatile yield. This could be visually conceived as checking that a SWRY trace line passes through a CPD defined QoI rectangle.

A few numerical methods are necessary for executing this analysis. SRWY model

temporal traces are evolved in time using an adaptive Huen-Euler integration scheme. The adaptive method was found to handle the system's stiffness while remaining computationally efficient. On a logarithmic time-scale between 1E-5 and 10 seconds, 200 points are saved for each volatile yield trace. Values representing the time to half the ultimate volatile yield are then located using a second-order Newton polynomial interpolation. This polynomial was chosen due to the curved nature of the volatile yield traces in the desired temporal periods. For random number generation the random.random function from the Python Numpy library was utilized to generate random samples from a continuos uniform distribution [67].

7.7 Consistency Analysis

From visualizing a sample of consistent points in Figure 7.13, it was evident that exploring a transformed parameter space would increase search efficiency. A simple linear transformation base on the apparent linear correlation between E and $\log_{10}(A)$ values was utilized: $\log_{10}(A) = \text{slope} \cdot E + \text{intercept}$. Instead of exploring the parameter spaces of A and E, the space of E and the intercept can be explored once a slope was fitted to consistent points found with an initial search. The greater efficiency of exploring the transformed space is demonstrated in Figure 7.13. The dash-dotted rectangular region represents the

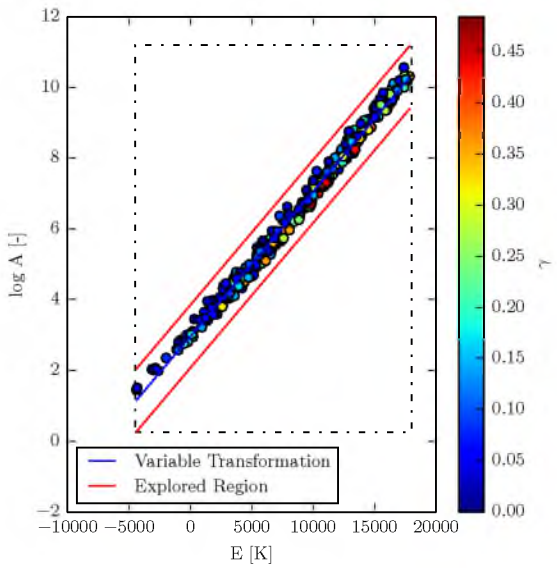

Figure 7.13: Visualization of variable transformation utilized for efficient exploration of parameter space. The dots represent consistent points in the parameter space and the dots' color indicates γ values. The variable transformation is shown through the blue line and the red lines indicate the region explored.

pure parameter hypercube that would be explored if a variable transformation was not used. The region between the two red lines is the transformed space that can be explored more efficiently. Throughout the process of refining the consistency search and regions of exploration, the transformation continuously evolved based upon the accumulation of additional data. A more refined search of the transformed E versus $\log_{10}(A)$ space is shown in Figure 7.14.

Linear transformed spaces for a and σ_a, as shown in Figure 7.15, were also utilized to minimize computational costs. Although initially a appeared to benefit from a variable transform, ultimately this transformation was found to be unnecessary and the nontransformed parameter space was utilized. The consistent points shown in Figure 7.14-7.15 were generated by taking 50,000 random samples from a four-dimensional transformed parameter space. Of those random samples, 1,244 consistent parameter-sets were found or approximately 2.5%.

By considering Figure 7.14-7.15 the consistent space can be noted to have many interesting characteristics. The spread of the consistent space is relatively narrow across many two-dimensional visualizations or marginals of the data, demonstrating that the utility of variable transformations may encompass more than just an increase in search efficiency. Bounds appear to exist for many of the parameters and shift throughout the four-dimensional space. For instance, parameter a seems limited to a range between 0.5 and 0.69, but the lower bound shifts upwards for the extreme E values.

Parameters E, A, and σ_a appear to be highly correlated, as could be deduced by the variable transforms used to explore their spaces. Such correlation was expected with fast heating rates because the coal was effectively experiencing a fixed temperature for a large temporal portion of the devolatilization. Looking back at Eq. (7.7), it can be shown that for a constant temperature the equation has nonunique solutions or correlation between three parameters:

$$\begin{aligned} K = \text{Constant} &= A_1 \exp\big(-(E_1 + \sigma_{a,1}Z)/T_P\big) \\ &= A_2 \exp\big(-(E_2 + \sigma_{a,2}Z)/T_P\big) \\ \ln(A_1) - \ln(A_2) &= \frac{E_1 + \sigma_{a,1}Z}{T_P} - \frac{E_2 + \sigma_{a,2}Z}{T_P}. \end{aligned} \tag{7.10}$$

With this strong correlation, the three parameters' bounds are also interdependent. The activation temperature E covered the entire parameter region explored, but tapered off in the number of consistent points found and respective γ values of those points in the limits of its explored region, -9,000 K to 40,000 K. Because of the nonunique behavior, it

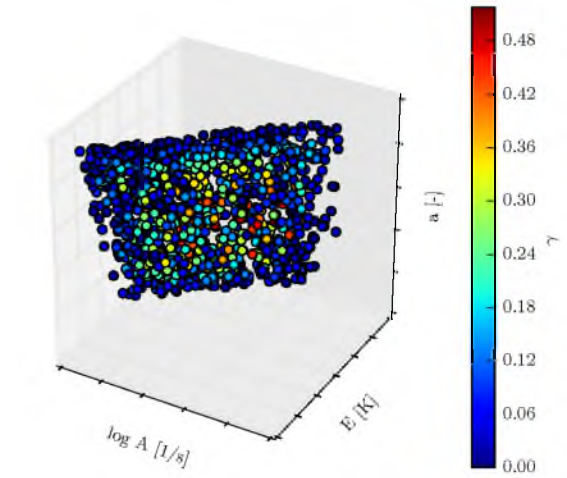

(a)

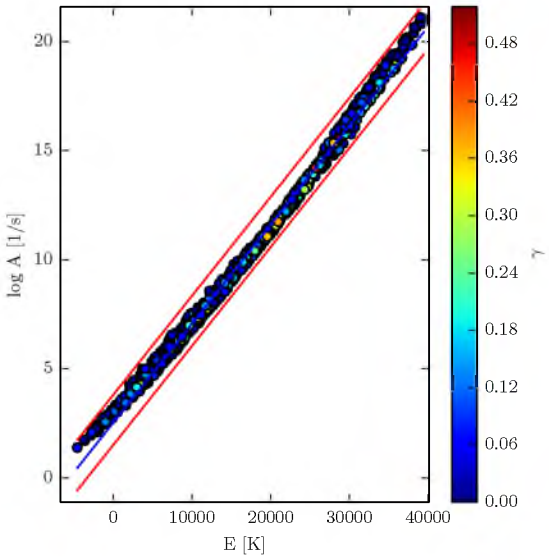

(b)

Figure 7.14: Marginal views of the SFWY model's four-dimensional parameter space. Dots signifying parameter-sets consistent with CPD uncertainty are shown in plot (a) across three parameter dimensions: $\log A$, E, and a. Plot (b) shows a two-dimensional view of the consistent points across $\log A$ and E space. The dot color corresponds to the respective γ value.

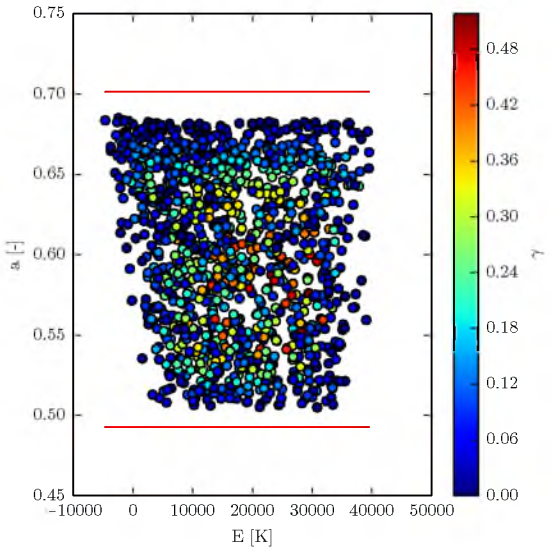

(a)

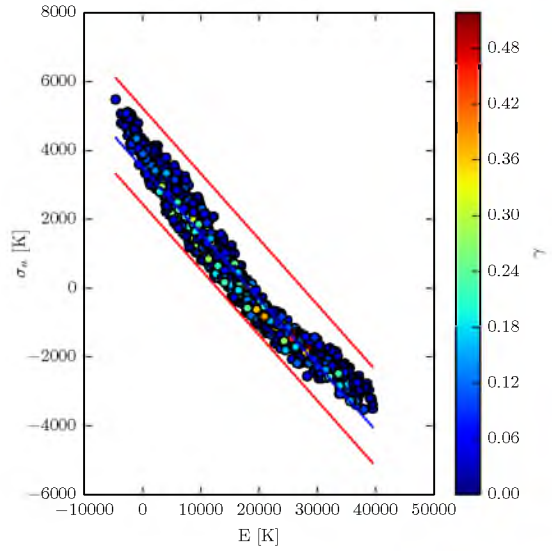

(b)

Figure 7.15: Visualization of two-dimensional parameter spaces explored during search of the four-dimensional parameter space. Plot (a) shows parameter a across parameter E and plot (b) likewise shows transformation of parameter σ_a with respect to E. Blue line indicates axis for transformed parameter space and red lines indicate region of parameter space explored.

is not possible to deduce which of the three correlated parameters limited the span of the consistent space.

Another interesting feature of the SRWY model found through the consistency search was the ability to remain consistent while inverting the temperature distribution. In Figure 7.15 it can be noted that consistent σ_a values become negative when the activation temperature E surpasses approximately 18,000 K. The negative sign should not be thought of as part of the σ_a value because a negative standard deviation is not possible. Instead, looking back at Eq. (7.7) and Figure 7.11, a negative value is indicative of inverting the activation temperature distribution, causing high activation temperatures to initially control the reaction and low activation temperatures to be operating when high amounts of conversion have occurred. This change in the temperature distribution should significantly alter the kinetic trends, yet evidently does so in a manner that maintains consistency. A visual comparison of the effect of inverting the temperature distribution can be seen in Figure 7.16, where two traces generated with consistent SRWY parameter-sets, but with opposite activation temperature distribution orientations, are compared with an equivalent nominal CPD trace.

It is evident that the original orientation of the temperature distribution produces a trace that bears characteristics more similar to that of the CPD trace. The trace from the inverted temperature distribution has a sharper slope, but does level-off at the desired volatile

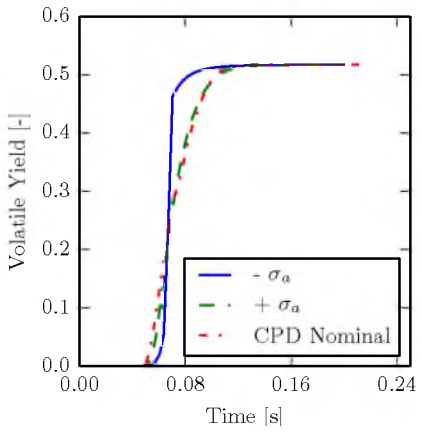

Figure 7.16: Comparing volatile yield traces of CPD and SRWY model with a hold-temperature of 2,400 K and heating-rate of 1E4 K/s. The CPD trace was calculated with nominal parameter values, while the SRWY traces were calculated with consistent parameter-sets. One SRWY trace utilized a consistent parameter set with the standard temperature distribution ($+\sigma_a$) and the other used an inverted temperature distribution ($-\sigma_a$).

yield. While consistent parameter-sets with inverted temperature distributions could be discounted from further consideration due to the poor shape characteristics, they will be retained presently due to fulfilling the current consistency criteria. If additional or redefined QoIs were utilized in the future, points with negative activation temperature distributions could justly be removed. This issue illustrates the distinction between adequacy and credibility. While the current consistency QoIs define model adequacy, they do not ensure credible solutions.

Throughout visualizations of consistent parameter-sets (Figure 7.13-7.15), it is evident that γ values associated with those consistent points are not continuously distributed across the four-dimensional space. It appears that parameter-sets with higher γ values typically lie within interior regions of the parameter space and that the bounding regions of the consistent parameter space have low gamma values, as would be expected. Exactly how the γ values are distributed across each of the four uncertainty SRWY parameters can be visualized within Figure 7.17.

All four uncertain parameters appear to have reasonably well defined distributions indicating that the parameter space exploration was sufficient. Truncation of the edges

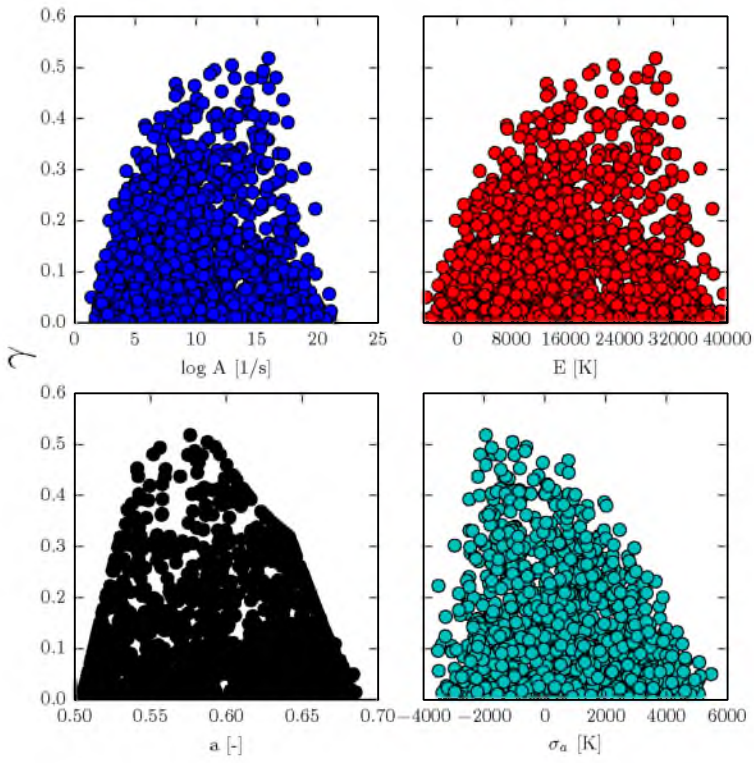

Figure 7.17: Scatter plots showing how γ values are distributed across the four free, SRWY model parameters for consistent parameter-sets.

of the distributions occurs, but should not significantly alter the distribution's appearance. The high-temperature ultimate volatile yield a correlates to γ with a distribution that reaches a maximum between an a value of 0.55 and 0.6. This is not surprising considering CPD predicted a similar range, as was seen in Figure 7.10. This distribution also has clear bounds that appear almost smooth even with the limited sample points. The most consistent parameter set or that with the highest γ value found corresponded to $\log_{10}(A) = 15.980\, s^{-1}$, $E = 29,400\, K$, $\sigma_a = -1,950\, K$, and $a = 0.576$. This set does not represent the largest γ possible for the SWRY model due to the use of random sampling, but acts as an estimate of the region of highest consistency. If consistent points with inverted temperature distributions are discounted, the most consistent set would be $\log_{10}(A) = 8.331\, s^{-1}$, $E = 13,240\, K$, $\sigma_a = 767.6\, K$, and $a = 0.580$.

Now that a set of consistent points has been located, the RPM's performance at representing the CPD model can be further evaluated. First, the consistent parameter-sets can be mapped into QoI space as shown in Figure 7.18. Clearly the red boxes, which were used as the QoI bounds within the consistency test, do not act as perfect representations of CPD's uncertainty, as is shown with the black dots taken directly from CPD's uncertainty investigation. With that caveat in mind, the red boxes do represent ranges of uncertainty believed to suffice in capturing the behaviors required of the SRWY model as specified by the application's requirements. Again, information incorporated into the QoIs outside of CPD's uncertainty is evident in comparing the red boxes and black dots.

Within Figure 7.18 the SRWY parameter a is shown as the color of the consistent points, or dots of colors other than black. Interestingly, there is no observable correlation to a values with the consistent point's positions within the 700 K DOE QoI spaces, but a linear correlation vertically across ultimate volatile yield is present in all other DOE points. The shape of the CPD uncertainty region at 700 K is also not captured. The rectangular shape of CPD's QoI uncertainty is well replicated by the SRWY model consistent points for all higher temperatures.

A few additional features of the QoI spaces are captured poorly by the SRWY model. For DOE points at 1,000 K, SRWY was not able to be consistent for ultimate volatile yields below 0.45, while CPD yields spanned to near 0.4. Similar issues are seen to lesser extents for higher temperature and lower heating-rate DOE points. Back in Section 7.4.4 it was decided that the SRWY would not include the heating-rate as a functional input. A consequence of this engineering decision can now be seen in the SRWY model's ability to match the QoIs across the range of heating-rates for all temperatures above 1,000 K.

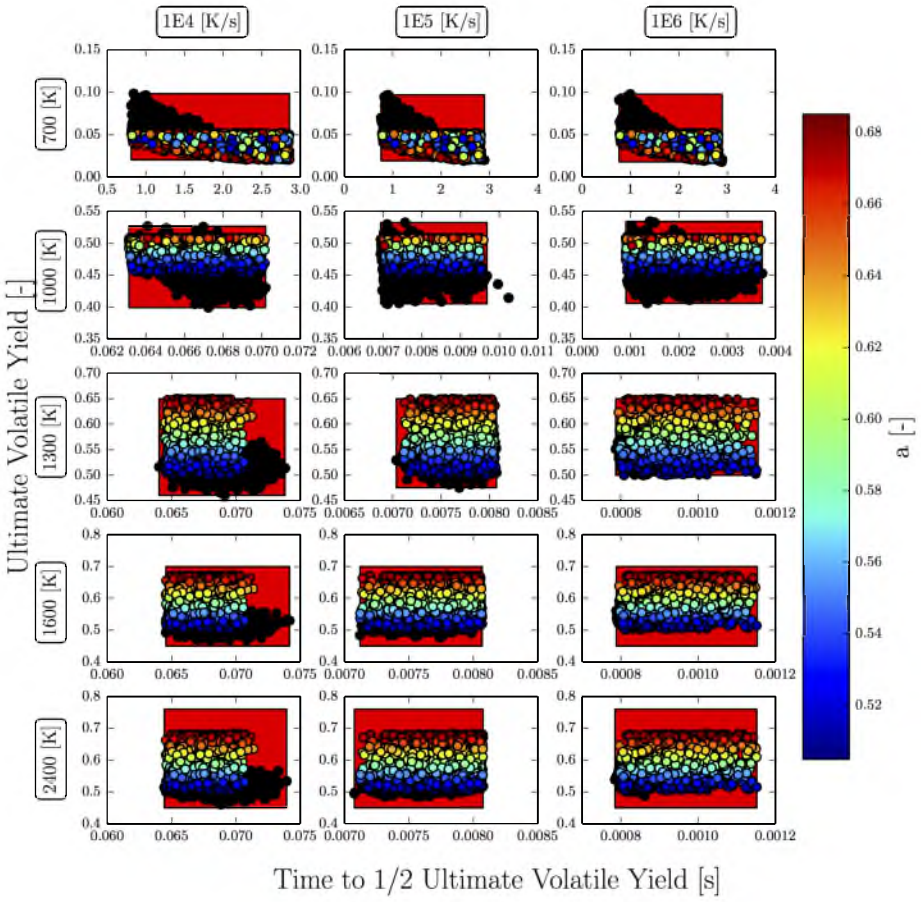

Figure 7.18: Comparison of uncertainty in QoI quantities from CPD calculations with equivalent values from RPM consistent parameter-sets across fifteen DOE conditions. Black dots indicate points generated from CPD uncertainty analysis, red boxes correspond to QoI spaces used to judge consistency, colored dots are consistent SRWY parameter-sets, and the color of the colored dots indicates the value of the high-temperature ultimate yield a for the consistent set. The plot's x axes are the time to half the ultimate volatile yield and y axes are the ultimate volatile yield value.

Consistent points are found in the initial portion of the temporal QoI's span for the 1E4 K/s DOE points, completely span 1E5 K/s DOE points, and overshoot much of CPD's uncertainty time-frame for 1E6 K/s DOE points in order to allow the model to span greater amounts of the temporal QoIs within other DOE conditions. As was anticipated the SWRY model performed best for the middle of the heating-rate range due to the yield model's parameter fitting method. The fact that QoI uncertainty ranges were on the same order of magnitude as the heating rate's effect upon the ultimate volatile yield allowed the SWRY model to find consistency, but the need to increase the temporal bounds for the 1E6 K/s high-temperature DOE conditions demonstrates the danger of this approximation.

7.8 Model Credibility

Up to this point the consistent parameter-sets for the RPM proved basic adequacy conditions had been met. Although consistent parameter-sets enable the SRWY model to meet specified QoI requirements, further analysis of the model's results is required to judge the credibility of the model-form for use beyond this application. While the model possesses characteristics specified by the QoIs, does it appear physical in other attributes not considered with the QoIs? Visualizing how the SRWY traces compare with CPD traces allows further comparison of characteristics not quantified with the current QoI definitions, as was previously shown in Figure 7.16, and is a fundamental view of reaction model performance. Analysis of model-form error or discrepancy between the RPM and CPD traces can provide evidence for the continued evolution of the model-form towards validity as well as qualified measures of credibility for the current model-form due to the credibility CPD possesses.

A comparison of 12 randomly selected consistent SRWY traces and CDP traces is shown in Figure 7.19. Only SRWY traces corresponding to noninverted activation temperature distributions are visualized due to the inverted distributions previously being noted to have poor shape characteristics. Kinetically, the SRWY traces begin to display substantial amounts of devolatilization within the same time-frame as the CPD traces and all traces match the initial kinetic shapes satisfactorily. Significant discrepancy then appears to occur within the two lower temperature DOE conditions during the later kinetic stages. As could also be noted in Figure 7.18, the SRWY traces do not reach as high of ultimate volatile yields as CPD traces at 700 K. For the DOE points with 1000 K hold-temperature and 1E4-1E5 K/s heating-rates, there is significant overshoot of the asymptote towards the ultimate volatile yield. A systematic discrepancy can be noted in how the traces corresponding to 1E4-1E5 K/s heating-rates asymptote to the ultimate volatile yield. The CPD traces have slower/gradual asymptotes, while the SRWY model's traces have sharper/abrupt asymptotes. Greater variance in the SRWY traces is also evident. Such variance reflects that the current QoIs are not over-constraining the SRWY model. The significant variation in the activation temperature distribution's standard deviation σ_a spanning from zero to approximately 10 percent of the activation temperature was expected to produce a wide assortment of kinetic shapes. Such variation due to the activation temperature distribution was most evident in the 1E6 K/s traces.

Following Ferson's (2008) [35] use of comparing cumulative distributions of QoIs, Figure 7.20 was constructed to visualize the comparison of uncertainties contained within

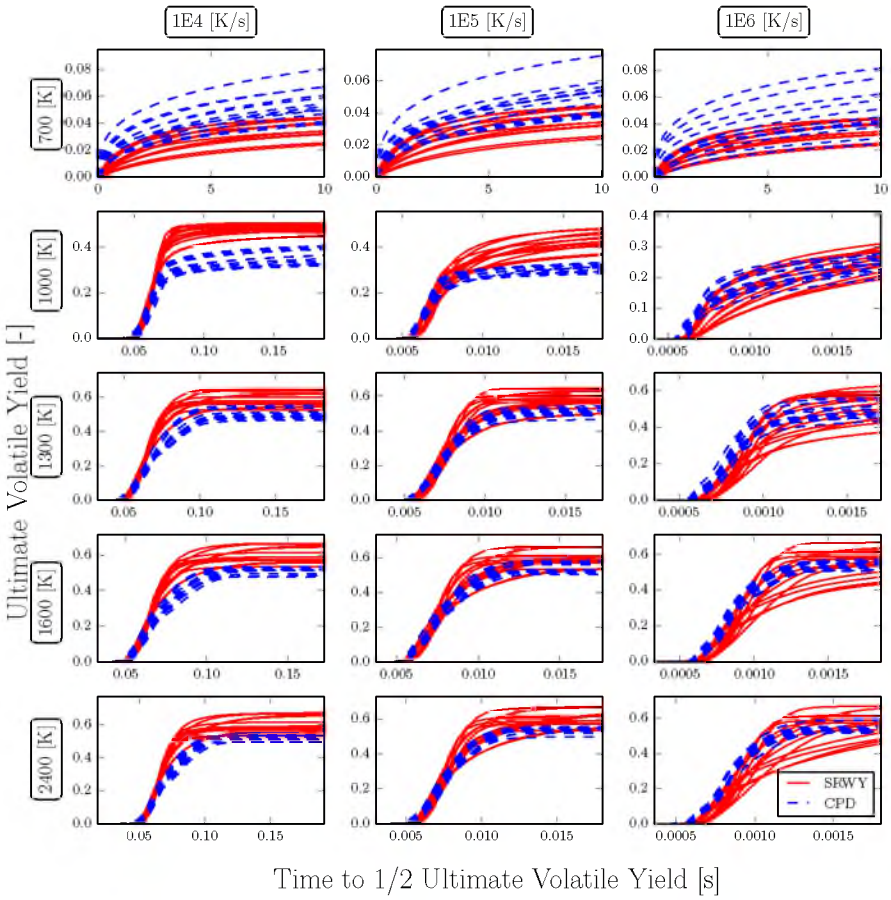

Figure 7.19: Comparing traces generated with 12 randomly selected consistent SRWY parameter-sets with 12 CPD traces created by randomly sampling the 13 uncertain parameters across 15 DOE conditions. Only SRWY parameter-sets with a noninverted temperature distribution are used in this comparison.

CPD and the RPM. These distributions should be utilized to visualize the RPM's ability to represent CPD's uncertainty characteristics, but the SWRY model's distributions should not be treated as real distributions as thought of in a probabilistic sense because they were produced by enforcing constraints. CPD's distributions are physical distributions and illustrate real propagation of uncertainty through the model into QoIs. All distributions shown are marginal distributions where one of the QoI dimensions was integrated out.

The most notable agreement between the SWRY model and CPD distributions can be found in the temporal QoIs for DOE points that used a 1E5 K/s heating-rate. Again the discrepancies previously noted can be found within the CDFs as well, with CPD's yield CDFs spanning larger ranges of yield QoI spaces for lower temperatures. With this visual testing methodology it is difficult to compare shapes for the higher temperature DOE points

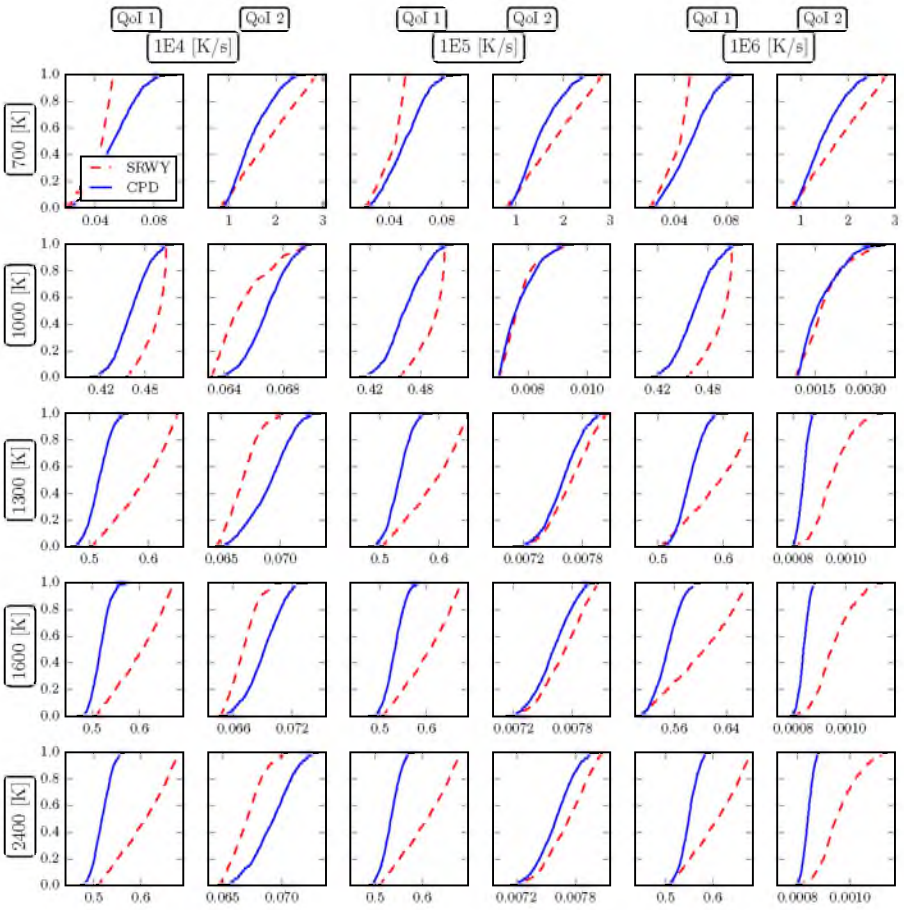

Figure 7.20: Comparison of quantities of interest's cumulative distributions for 15 design of experiment conditions. Distributions were marginalized over other QoI, respectively. The blue lines represent CPD and the dashed red lines stand for the SWRY model's consistent points.

due to the expanded QoI bounds utilized, which distort the distributions and to which CPD is not expected to be directly comparable. Ideally, the more the SWRY model's CDFs resemble CPD's, the better the forward propagating the SWRY model's uncertainty through the application would represent the forward propagation of CPD's uncertainty.

Overall, the SRWY model's performance was deemed satisfactory for the application requirements. While shortcomings of the SRWY model have been highlighted, its ability to meet strenuous requirements for consistency with CPD characteristics across a wide range of system conditions is a strong statement towards its credibility. The traces produced by the SWRY model closely resemble the equivalent CPD traces and its ability to represent CPD's propagated uncertainty has been qualified through CDFs. Basing our evaluation of the SWRY model on the application space demands for a scale-bridging approximation

of CPD that could capture thermodynamic yield trends, the SWRY model was deemed to meet the requirements.

7.9 Model Refinement

Although time constraints necessitated the use of the model-form described thus far for the first year's simulations, utilizing the discrepancies observed to motivate model-form refinement for year two's simulations was possible. Propagation of knowledge gained is a fundamental piece of the development of credible models. Due to the discrepancy correlated to low temperatures, factors effecting the model's performance in this condition region were reconsidered.

Within Section 7.4.1 a hold-time of 10 seconds was assumed to be the effective equilibrium time-scale for the application. Figure 7.4 demonstrated that this assumption was not ideal for temperatures below 1,200 K even when the coal effectively experienced instantaneous heating. While it is true that the coal will spend less than 10 seconds in the boilers, this assumption reduced the driving force for reactions at lower temperatures. The consequences of a reduced driving force were carried though the model development process, effecting the consistent spaces for the free parameters explored. This is an illustrative example of the difficulty of performing calibration and quantifying model-form uncertainty simultaneously, also known as the identification problem [6]. More realistic estimates of the time to reach equilibrium were investigated within Figure 7.21, where nominal CPD parameter values were again utilized. Estimates of time to equilibrium were found by locating positions on yield traces where $(V_f - V)/V_f < 0.01$, or the yield was within one percent of the final yield for that hold-temperature.

As was previously noted, 10 seconds is a good approximation of an equilibrium hold-time above 1200 K, but now it becomes evident that this is a poor assumption for the lower temperatures included in the DOE. Data-points for temperatures below 900 K were not included in Figure 7.21 because the time to compute such data was prohibitive and the scaling behavior had already become evident. Another important consideration is that this figure was created using a 1E9 K/s linear heat-up rate. Slower heat-up rates would most strongly effect lower temperatures, lengthening the time to reach equilibrium. Although the creation of an accurate yield model for low temperatures is out of the scope of this study, increasing the hold-times for 500, 600, 700, 800 and 900 K CPD runs to 1E5 seconds and 1000 K to 1E4 seconds was viable. The effect of incorporating this data into an improved yield model can then be utilized to determine if future model development iterations should

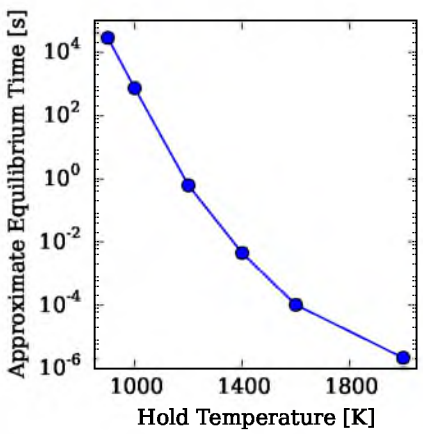

Figure 7.21: Approximate thermodynamic equilibrium time for CPD traces across a range of hold-temperatures. CPD traces were generated using nominal parameter values and a linear heating-rate of 1E9 K/s from 300 K to the respective hold-temperature.

focus upon this issue.

Using the same methodology as was previously described within Section 7.5, a yield model was created from CPD data that spanned a range of heating-rates and hold-temperatures, where the hold-times for lower temperatures were increased. This adjusted yield model is shown in Figure 7.22. The parameters fit to the CPD data were b = 11.53, c = -9.122, d = 2.407, e = -0.7773, and 500 K was the devolatilization initiation temperature. The effect of longer hold-times becomes evident when the adjusted yield model is compared with Figure 7.10. Also included in Figure 7.22 are purple dashed-lines indicating potential model-forms that could be explored in the future to account for uncertainty in the yields at higher temperatures. Until data are available in such temperature regions, it will be difficult to compare potential model-forms and reach definitive conclusions. Even if high-temperature experimental data for pure devolatilization do not become available in the near future, model-forms such as those suggested could be tested within multiphysics simulations against data-forms available for comparison at that scale, or a top-down validation.

The updated yield model was utilized in Eq. (7.7) and consistent parameter-sets were located in the manner previously described within Section 7.6. Plots of the consistent parameter-sets can be found in Appendix A, but the effect of the improved yield model can be judge through a yield trace comparison, as shown in Figure 7.23. Compared to Figure 7.19, improved performance in matching the low-temperature yield trends is evident. Additionally, the traces demonstrated improved matching of CPD's asymptotic behavior across the higher temperature traces. The improvement in the asymptotic behavior appears

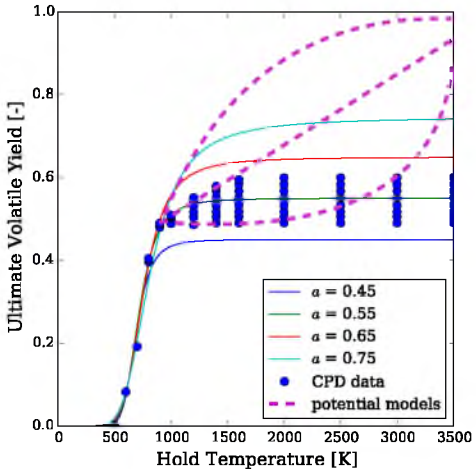

Figure 7.22: Yield model equivalent to Figure 7.10 once the hold-times were increased for lower temperatures. Lines are based upon Eq. (7.3), where each line signifies a different a value and the dots are CPD data with the heating-rates and hold-temperatures varied. Additional purple dashed lines represent potential alternative forms of the yield model that could be implemented to account for uncertainty at higher temperatures.

to be at least in part caused by the activation temperature distribution's standard deviation σ_a now varying between 3,000 and 8,000 K, where previously it ranged from 0 to 5,000 K. The distribution's width is of similar breadth, but the updated yield model caused a shift towards larger values. The variance previously noted in the appearance of the volatile yield trends has been significantly reduced. While again this is likely due to a combination of factors, the large reduction in widths of the consistent ranges for A and E are likely major contributors. With the updated yield model, optimal parameter-set values were found to be $\log_{10}(A) = 8.499\,s^{-1}$, $E = 14,380\,K$, $\sigma_a = 4,719\,K$, and $a = 0.565$. These optimal values are similar to those previously found, when inverted energy distributions were not considered, except that the activation temperature distribution's value is approximately six times larger.

An additional attribute of applying the SWRY model on the current application's DOE becomes evident once the implications of Figure 7.21 for high temperatures is considered within Figure 7.23. For any given heating-rate, the traces above 1,000 K appear to be approximately the same. For the fastest heating-rate considered, 1E6 K/s, it takes approximately 0.001 seconds to reach 1,000 K and the thermodynamic time-scale is over 100 seconds at that temperature. However, for 1600 K the thermodynamic time-scale is approximately 1E-4 seconds, meaning that the traces are heating-rate limited and effectively are thermodynamic yield curves. Evidently, future iterations of model-development could

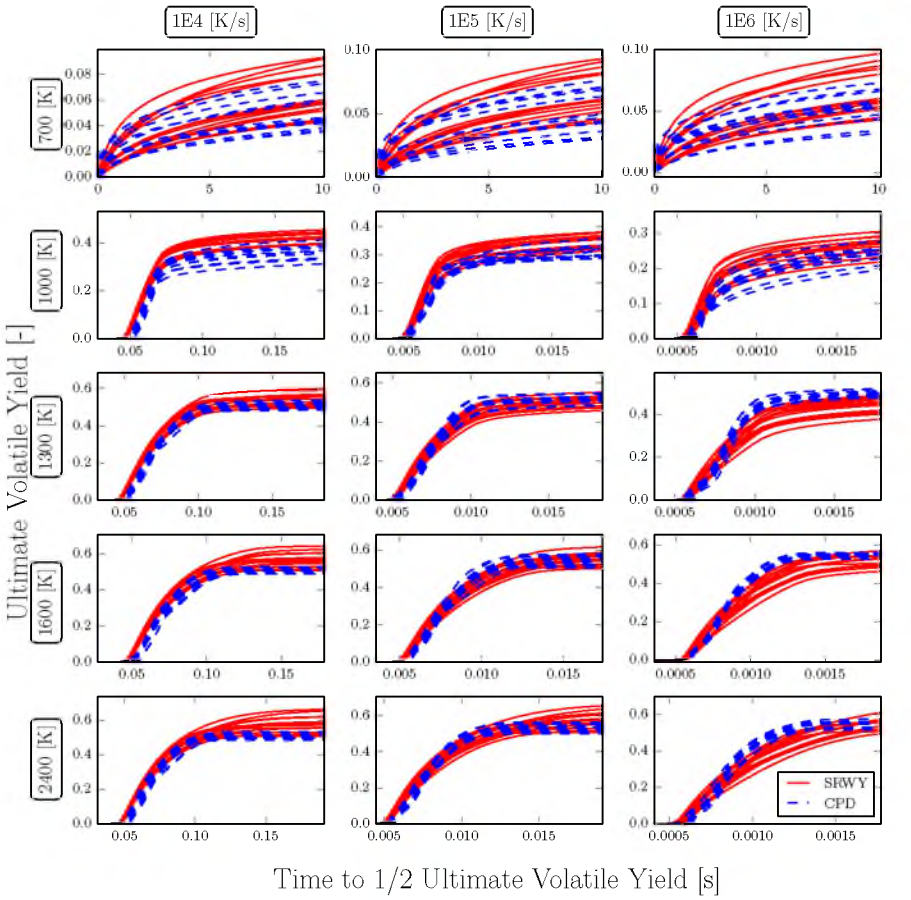

Figure 7.23: A comparison of 12 traces generated with randomly selected consistent SRWY parameter-sets with 12 CPD traces created by randomly sampling the 13 uncertain parameters across 15 DOE conditions. The SRWY model used the yield model-form based upon longer times to reach equilibrium for lower temperatures.

remove the 2,400 K DOE conditions and likely benefit from additional lower temperature DOE points or higher heating-rate points where the heating-rate conditions could be isolated from the kinetics. Moving forward it can also be noted within Appendix A that the extend time range QoI for higher temperatures and 1E6 K/s heating rates could be removed.

Clearly the updated yield model positively effected the RPM's credibility. Given the performance noted across traces for all DOE conditions, there should be greater confidence in utilizing the RPM in extrapolating applications. Throughout the developmental process of creating the current model-form, tasks that could be completed to increase the model's validity for lower or higher temperature applications have been detected. Lower temperature applications would benefit from additional refinement of the yield model through longer CPD runs. Higher temperature uses could look to the creation of experimental data,

or inversely, validate through comparing multiphysics simulations with experimental data available at that scale.

7.10 Conclusions

The need for a computationally cheap function to capture desired traits of a rigorous physics model, which has been deemed to contain significant amounts of uncertainty, was the driving force behind this research. Once the desired characteristics of the more rigorous model and corresponding uncertain had been quantified, a reduced physics based model was created using a single reaction model with yield model and distributed activation energy. This RPM contained four free parameters, which were calibrated using a consistency test against selected QoIs. The QoIs were based upon capturing desired physics and were quantified by the uncertainty contained within the more rigorous model and additional insights. Once consistent parameter-sets were located, the credibility of the reduced physics model in representing the more rigorous model was evaluated. Through visualizing the mapping of the consistent parameter-sets into the space of the QoIs, qualitatively comparing the characteristics of kinetic traces and considering cumulative distributions, the RPM was deemed to have satisfactory performance for the stated application. Then utilizing the discrepancies discovered throughout the model validation, an improvement to the yield model was implemented and the gained performance demonstrated.

The model development and analysis demonstrated throughout this work were based upon one particular coal type, Utah Sufco bituminous. The SWRY model can be applied to alternative coal types simply through repeating steps utilized during the model development process. The most significant difference between coal types will come from fitting the yield model to coal specific CPD thermodynamic data. Once alternative fitting parameters have been found for the yield model, the consistency test can be rerun. The parameter space exploration should be expedited by utilizing the consistent regions found for the current coal type as prior knowledge to base parameter bounds upon.

Like all engineering exercises, this process has the potential for continued refinement. Incorporation of additional sources of data, especially experimental data, is an obvious next step. High-temperature experimental-data could greatly reduce the uncertainty in the model-form for higher-temperature applications. An alternative route to experimental data could be to compare against another high-fidelity model such as FLASHCHAIN [101], which could give further credibility to the RPM within its current areas of application. Continued alteration to the form of the yield model is another avenue of simple improvement

that has been demonstrated to positively impact the model's performance and which could include model-form comparisons. Reconsidering the DOE design is another easy alteration to incorporate in order to gain additional insight efficiently. As was previously stated, QoI definitions are subjective, so further exploration might lead to superior model performance. Balancing the cost-benefit of such refinements is a research area in its own right.

CHAPTER 8

CONCLUSIONS

The narrative this dissertation aimed to espouse has been the demonstration of how scale-bridging models are created and how adoption of validation and uncertainty quantification methods into the development process for scale-bridging models generates greater model credibility. Scale-bridging models are an unusual type of model that are not meant to describe a novel phenomena nor capture a phenomena to a higher degree of accuracy than previous models of the same phenomena. Instead scale-bridging models are ultimately an engineering based approach towards creating the appropriate tool to meet application mandates. Simply fitting a lower-fidelity model to desired model outputs of a higher-fidelity model can create a functional scale-bridging model, but inevitably it will be desired to use that model outside of the range of data to which it was fit. To have confidence in the scale-bridging model's robustness, greater understanding of the model's development process, sensitivities, limitations, and strengths is necessary. Those desired attributes can be described as the model's credibility. Methodologies from the validation and uncertainty quantification community provide means of increasing the credibility of scale-bridging models. An overview of material presented, accomplishments, and ideas put forth will now be summarized.

A novel modeling framework for capturing multiple-polymorph precipitation in highly supersaturated, aqueous conditions was formulated. The novelty of this framework was in its application system, use of moment methods, and detailed physics suite. This framework was validated for a $CaCO_3$ system, demonstrating the framework's strengths and limitations. The framework and validation thereof were published within Schroeder et al. (2014) [127]. This framework was developed as a scale-bridging model and its adoption into computational fluid dynamics codes was demonstrated within Abboud et al. (2015) [1].

An overview of two validation and uncertainty quantification approaches were outlined and these methods were demonstrated further in subsequent material. A philosophical approach describing a cyclic process for developing credible simulations was then presented.

This philosophy was similar to methods previously presented, but its focus upon the conceptual model, as well as its justification and qualification processes were unique. This approach was then demonstrated on a pedagogical example that can be utilized as future educational material.

Model-form uncertainty refers to the effort to quantify the amount of error within model predictions attributed to the imperfections in model-form. A Bayesian framework known as the Kennedy O'Hagan approach was explored through the application to a chemical kinetics pedagogical example. This application walkthrough illustrated strengths and limitations of the current methodology. Additional ideas within this realm were also explored such as multi-input and constrained Gaussian Processes.

A novel scale-bridging model for capturing coal particle enthalpy and heat capacity was presented. This model was created to meet a specific application demand of reducing a relatively expensive iterative solve. Through employment of Bayesian parameter estimation and Markov chain Monte Carlo methods, piecewise model-forms were created and calibrated to meet application requirements. While the development process created an adequate product, the process itself shed light on the uncertainty within this piece of physics that was previously not explicitly acknowledged.

Coal devolatilization was another piece of coal physics determined to need scale-bridging. A novel model-form was developed and consistency constraints were applied to locate consistent parameter regions. The uncertainty contained within a high-fidelity devolatilization model was used to define error bounds for the consistency analysis. After consistent parameter regions were located, subsequent analysis determined that the model was adequate for the application. Mapping consistent parameter sets into the space of the constraints and comparing trace shapes across the conditions specified by the application-based design of experiments were key tools for the validation assessment. Further consideration of the discrepancies within the scale-bridging model's performance led to model refinements that improved the model's performance and increased credibility. Future areas of emphasis for continued model development were also noted throughout the analysis of the current model-form.

APPENDIX A

ADDITIONAL FIGURES FOR CHAPTER 7

A.1 Figures for Refined Devolatilization Model

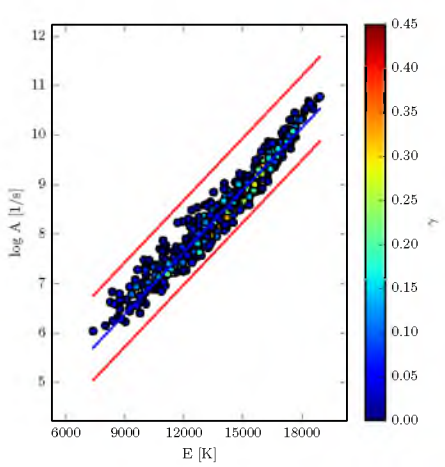

Figure A.1: Two-dimensional view of the consistent points across $\log A$ and E space once the yield model was improved as specified within Section 7.9. The dot color corresponds to the respective γ value. The variable transformation utilized to increase search efficiency is shown through the blue line and the red lines indicate the region explored.

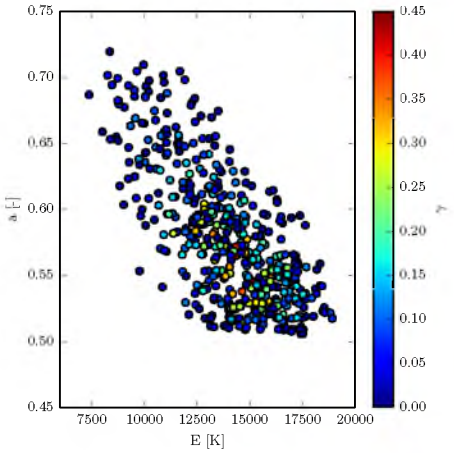

Figure A.2: Two-dimensional view of the consistent points across a and E space once the yield model was improved as specified within Section 7.9. The dot color corresponds to the respective γ value.

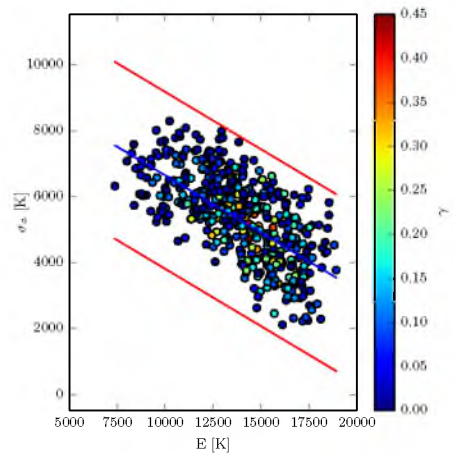

Figure A.3: Two-dimensional view of the consistent points across σ_a and E space once the yield model was improved as specified within Section 7.9. The dot color corresponds to the respective γ value. The variable transformation utilized to increase search efficiency is shown through the blue line and the red lines indicate the region explored.

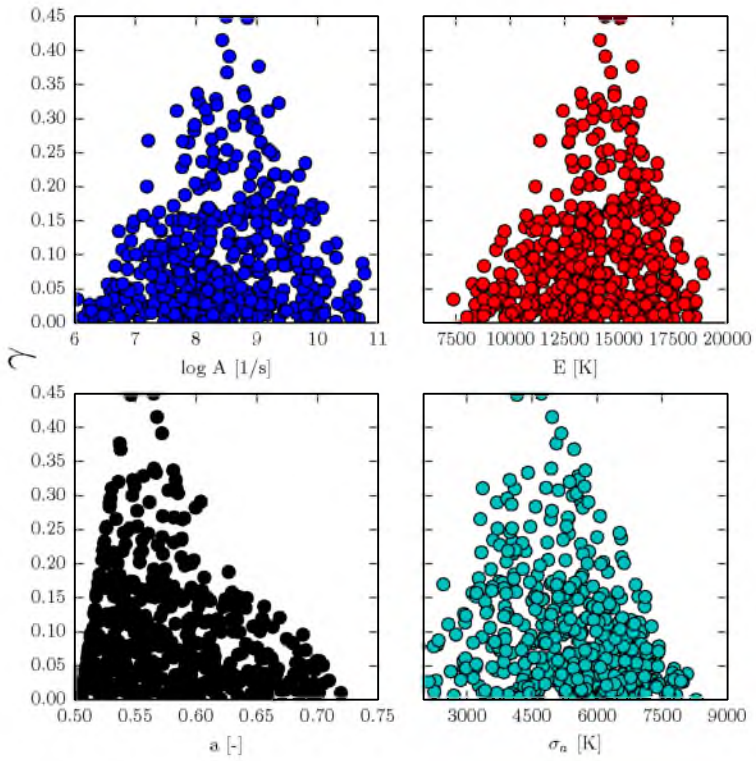

Figure A.4: Scatter plots showing how γ values are distributed across the four free SRWY model parameters for consistent parameter sets once the yield model was updated in the manner suggested within Section 7.9.

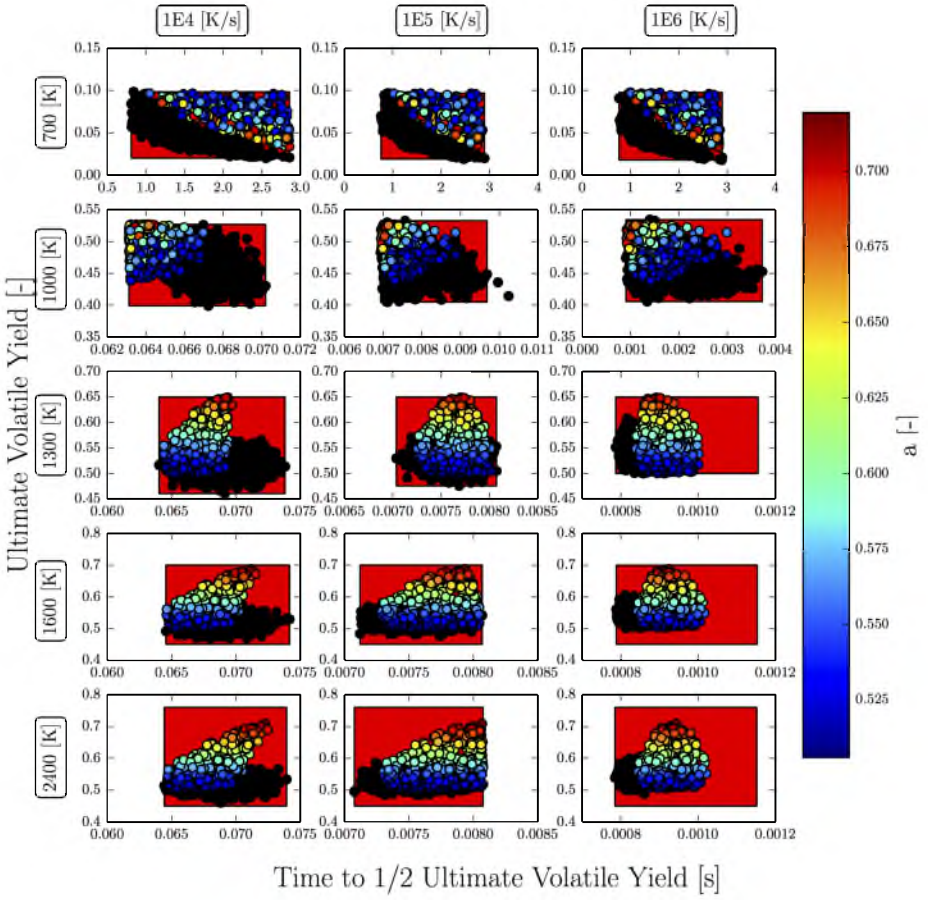

Figure A.5: Comparison of uncertainty in QoI quantities from CPD calculations with equivalent values from RPM consistent parameter sets across 15 DOE conditions. The RPM included the improved yield model, as specified within Section 7.9. Black dots indicate points generated from CPD uncertainty analysis, red boxes correspond to QoI limits used in judging consistency, colored dots are consistent SRWY parameter sets, and the color of the colored dots indicate the value of the high temperature ultimate yield value a for the consistent set. Plot's x axis is the time to half the ultimate volatile yield and y axis is the ultimate volatile yield value.

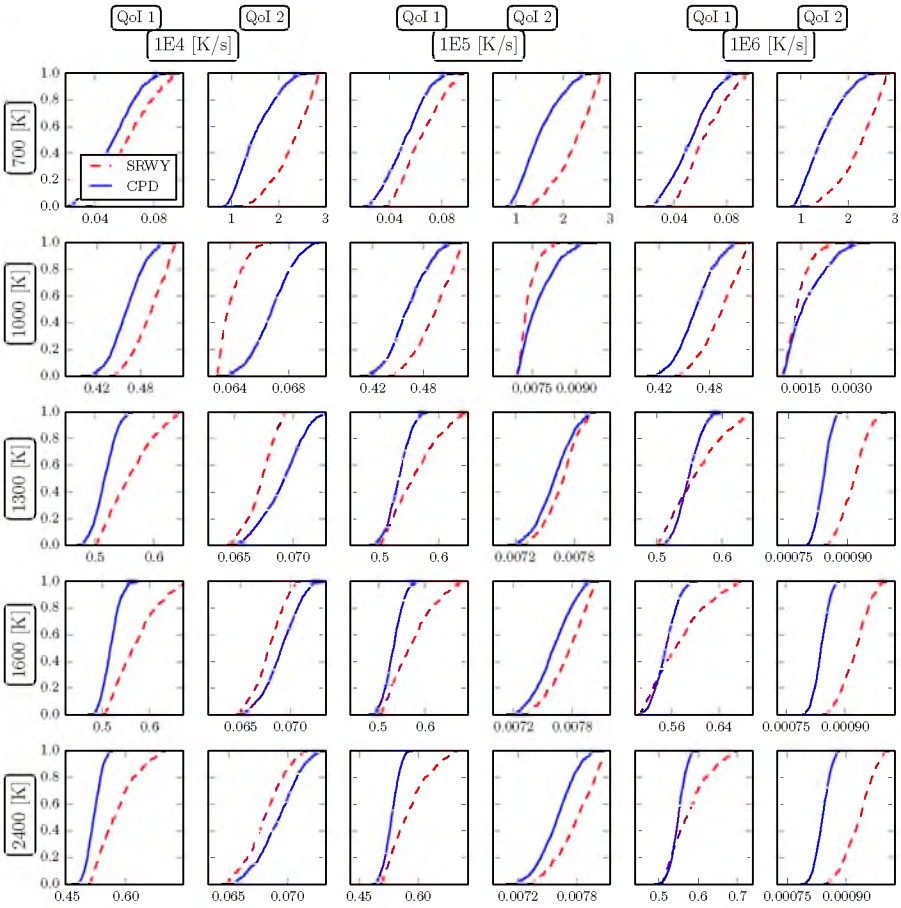

Figure A.6: Comparison of quantities of interest's cumulative distributions for 15 design of experiment conditions. Distributions were marginalized over other QoI, respectively. The blue lines represent CPD and the dashed red lines stand for the SWRY model's consistent points. SWRY model utilized improved yield model suggested within Section 7.9.

APPENDIX B

HELPFUL TOOLS

B.1 Software Tools Utilized

The following list includes software tools found useful through the course of my graduate studies, even beyond the material presented throughout the dissertation.

- *Python*: open source programing language utilized with Numpy and Scipy libraries to perform majority of numerical calculations.
  `https://www.python.org`
- *pyDOE*: open source Python library used to create Latin Hypercube designs.
  `http://pythonhosted.org/pyDOE/`
- *emcee*: open source Python based Markov chain Monte Carlo tool utilized within Bayesian analysis to explore multidimensional posterior distributions.
  `http://dan.iel.fm/emcee/current/`
- *Spyder*: open source GUI interface for Python that was utilized for creating, running, and debugging many Python scripts.
  `https://github.com/spyder-ide/spyder`
- *pyregress*: in-house developed, Python based Gaussian process tool. For more information contact `sean.t.smith@utah.edu`
- *CPD*: MATLAB version of CPD code developed at Brigham Young University. For more information contact `tom_fletcher@byu.edu`
- *matlab.engine*: Python library for calling MATLAB functions from within Python scripts.
- *SourceTree*: Free Git repository management software.
  `https://www.sourcetreeapp.com`

- *cpd-matlab*: Git repository for scripts used for coal devolatilization work.
  `https://bitbucket.org/team_sean/cpd-matlab`

B.2 Coal Devolatilization Algorithm

Algorithmic steps followed for current devolatilization results. Scripts listed found in cpd-matlab Git repository.

- Collect coal specific parameter values and uncertainties for CPD. Potentially from Professor Fletcher at Brigham Young University.
- Create volatile yield traces over a range of heating-rates and hold-temperatures using MATLAB version of CPD. Export data for continued analysis within Python.
  Script: cpd_explore.m
- Fit CPD ultimate volatile yield data (spanning conditions from previous step) to desired form of yield equation such as Eq. (7.3). Can be fit with Scipy's minimize function using the Nelder-Mead algorithmic option.
  Script: model_form_fitting.py
- Determine quantities of interest (i.e., ultimate volatile yield) and system conditions of interest (i.e., heating-rates and hold-temperatures) that will be utilized for consistency tests.
- Define uncertainty within CPD for all QoIs and system conditions of interest through Monte Carlo sampling of uncertain parameters. These data are then used as consistency constraint bounds within consistency tests. Monte Carlo sampling can be completed using Numpy's random.random function and the MATLAB version of CPD can be called within Python scripts with the matlab.engine library.
  Script: uncertainty_explore.py
- Explore the reduced model's parameter space with a Monte Carlo search, where parameter correlations as shown in Section 7.7 increase search-efficiency. Iterate this step to improve number of consistent points found by improving search region. Again, Numpy's random.random function can be used for the random search. Save consistency search results in a format amenable to future analysis and plotting.
  Script: consistency_test.py
- Visualize consistent points in parameter space, QoI space, and through yield traces for continued analysis and utilize discrepancies to motivate continued model-form

evolution.

Scripts: consistency_visualize.py and compare_rpm_cpd.py

REFERENCES

[1] A. Abboud, B. Schroeder, T. Saad, S. Smith, D. Harris, and D. Lignell, *A numerical comparison of precipitating turbulent flows between large-eddy simulation and one-dimensional turbulence*, AIChE J., Accepted May (2015).

[2] J. Abrahamson, *Graphite sublimation temperatures, carbon arcs and crystallite erosion*, Carbon, 12 (1974), pp. 111–118.

[3] J.-P. Andreassen, *Formation mechanism and morphology in precipitation of vaterite–nano-aggregation or crystal growth?*, J. Cryst. Growth, 274 (2005), pp. 256–264.

[4] J.-P. Andreassen and M. J. Hounslow, *Growth and aggregation of vaterite in seeded-batch experiments*, AIChE J., 50 (2004), pp. 2772–2782.

[5] D. Anthony, J. Howard, H. Hottel, and H. Meissner, *Rapid devolatilization of pulverized coal*, 15th Symp. (Int) Combust., 1303 (1975).

[6] P. Arendt, D. Apley, and W. Chen, *Quantification of model uncertainty: Calibration, model discrepancy, and identifiability*, J. Mech. Des., 134 (2012), p. 100908.

[7] P. Arendt, D. Apley, W. Chen, D. Lamb, and D. Gorsich, *Improving identifiability in model calibration using multiple responses*, J. Mech. Des., 134 (2012), p. 100909.

[8] ASME, *Guide for verification and validation in computational solid mechanics*. Revised Draft, March 2006.

[9] S. Badzioch and P. Hawksley, *Kinetics of thermal decomposition of pulverized coal particles*, Ind. Eng. Chem. Process Des. Dev., 9 (1970), p. 521.

[10] J. Barthel, H. Krienke, and W. Kunz, *Physical Chemistry of Electrolyte Solutions: Modern Aspects*, vol. 5 of Topics in Physical Chemistry, Springer, Triltsch, Würzburg, 1998.

[11] M. Bayarri, J. Berger, R. Paulo, J. Sacks, J. Cafeo, J. Cavendish, C.-H. Lin, and J. Tu, *A framework for validation of computer models*, Tech. Rep. 162, National Institute of Statistical Sciences, 2005.

[12] ——, *A framework for validation of computer models*, Technometrics, 49 (2007), pp. 138–154.

[13] R. Beck and J.-P. Andreassen, *The onset of spherulitic growth in crystallization of calcium carbonate*, J. Cryst. Growth, 312 (2010), pp. 2226–2238.

[14] E. Biagini and L. Tognotti, *A generalized correlation for coal devolatilization kinetics at high temperatures*, Fuel Process. Technol., 126 (2014), pp. 513–520.

[15] P. Bots, L. Benning, J. Rodriguez-Blanco, T. Roncal-Herrero, and S. Shaw, *Mechanistic insights into the crystallization of amorphous calcium carbonate (ACC)*, Cryst. Growth Des., 12 (2012), pp. 3806–3814.

[16] G. Box and N. Draper, *Empirical Model-Building and Response Surfaces*, John Wiley & Sons, Inc., New York, 1987, p. 74.

[17] P. Boyle, *Gaussian Processes for Regression and Optimisation*, PhD thesis, Victoria University of Wellington, Victoria University of Wellington, 2007.

[18] J. Bracco, A. Stack, and C. Steefel, *Upscaling calcite growth rates from the mesoscale to the macroscale*, Env. Sci. Technol., 47 (2013), pp. 7555–7562.

[19] L. Brečević and A. Nielsen, *Solubility of amorphous calcium carbonate*, J. Cryst. Growth, 98 (1989), pp. 504–510.

[20] J. Brynjarsdottir and A. O'Hagan, *Learning about physical parameters: The importance of model discrepancy*, Inverse Probl., (2014), p. WEB Sept. 2014.

[21] J. Cartwrigth, A. Checa, J. Gale, D. Gebauer, and C. Sainz-Días, *Calcium carbonate polyamorphism and its role in biomineralization: How many amorphous calcium carbonates are there?* Angew. Chem., Int. Ed., 51 (2012), pp. 11960–11970.

[22] L. Chou, R. Garrels, and R. Wallast, *Comparative study of the kinetics and mechanisms of dissolution of carbonate minerals*, Chem. Geol., 78 (1989), pp. 269–282.

[23] J. Christoffersen and M. Christoffersen, *Kinetics of spiral growth of calcite crystals and determination of the absolute rate constant*, J. Cryst. Growth, 100 (1990), pp. 203–211.

[24] P. Cubillas, S. Köhler, M. Prieto, C. C., and E. Oelkers, *Experimental determination of the dissolution rates of calcite, aragonite, and bivalves*, Chem. Geol., 216 (2005), pp. 59–77.

[25] S. Da Veiga and A. Marrel, *Gaussian process modeling with inequality constraints*, Ann. Fac. Sci. Toulouse Math., 21 (2012), pp. 529–555.

[26] R. David, F. Espitalier, and A. Caeirão, *Developments in the understanding and modeling of the agglomeration of suspended crystals in crystallization from solutions*, KONA, 21 (2003), pp. 40–53.

[27] J. Dirksen and T. Ring, *Fundamentals of crystallization: Kinetic effects on particle size distributions and morphology*, Chem. Eng. Sci., 46 (1991), pp. 2389–2427.

[28] M. Donnet, P. Bowen, and J. Lemaître, *A thermodynamic solution model for calcium carbonate: Towards an understanding of multi-equilibria precipitation pathways*, J. Colloid Interface Sci., 340 (2009), pp. 218 – 224.

[29] E. Donskoi and D. McElwain, *Optimization of coal pyrolysis modeling*, Combust. Flame, 122 (2000), pp. 359–367.

[30] M. Elimelech, J. Gregory, X. Jia, and R. Williams, *Modeling of Aggregation Processes*, Colloid and Surface Engineering Series, Butterworth-Heinemann, New York, 1998, ch. 6.

[31] K.-T. Fang, R. Li, and A. Sudjianto, *Design and Modeling for Computer Experiments*, Computer Science & Data Analysis Series, Chapman & Hall/CRC, New York, 2005.

[32] R. Feeley, *Fighting the Curse of Dimensionality: A Method for Model Validation and Uncertainty Propagation for Complex Simulation Models*, PhD thesis, University of California, Berkeley, 2008.

[33] R. Feeley, M. Frenklach, M. Onsum, T. Russi, A. Arkin, and A. Packard, *Model discrimination using data collaboration*, J. Phys. Chem. A, 110 (2006), pp. 6803–6813.

[34] R. Feeley, P. Seiler, A. Packard, and M. Frenklach, *Consistency of a reaction dataset*, J. Phys. Chem. A, 108 (2004), pp. 9573–9583.

[35] S. Ferson, L. Oberkampf, and L. Ginzburg, *Model validation and predictive capability for the thermal challenge problem*, Comput. Method Appl. Mech. Engrg., 197 (2008), pp. 2408–2430.

[36] T. Fletcher and T. Holland, *CPD Heat.* `http://www.et.byu.edu/~tom/cpd/cpdcodes.html`, November 2014.

[37] T. Fletcher, A. Kerstein, R. Pugmire, M. Solum, and D. Grant, *A chemical percolation model for devolatilization: Summary*, Sandia Technical Report 92-8207, Sandia National Laboratories, 1992.

[38] D. Foreman-Mackey, D. Hogg, D. Lang, and J. Goodman, *emcee: The MCMC hammer*, PASP, 125 (2013), pp. 306–312.

[39] R. Fox, *Computational Model for Turbulent Reacting Flows*, Cambridge Series in Chemical Engineering, Cambridge University Press, New York, 2003.

[40] M. Frenklach, *Transforming data into knowledge–process informatics for combustion chemistry*, P. Combust. Inst., 31 (2007), pp. 125–140.

[41] M. Frenklach, A. Packard, and P. Seiler, *Prediction uncertainty from models and data*, P. Amer. Contr. Conf., (2002).

[42] D. Gebauer, A. Völkel, and H. Cölfen, *Stable prenucleation calcium carbonate clusters*, Science, 322 (2008), pp. 1819–22.

[43] T. Gebrehiwet, G. Redden, Y. Fujita, M. Beig, and R. Smith, *The effect of the CO_3^{2-} to Ca^{2+} ion activity ratio on calcite precipitation kinetics and Sr^{2+} partitioning*, Geochim. T., 13 (2012).

[44] A. Gelman, J. Carlin, H. Stern, D. Bunson, A. Vehtari, and D. Rubin, *Bayesian Data Analysis*, Chapman & Hall/CRC Texts In Statistical Science, Chapman and Hall/CRC, New York, 2013.

[45] J. Gibbs, *Collected Works*, vol. I, Longmans, Green and Co., New York, 1928.

[46] W. Gilks, S. Richardson, and D. Spiegelhalter, eds., *Markov Chain Monte Carlo in Practice*, Chapman & Hall/CRC, New York, 1996.

[47] S. Girshick and C. Chiu, *Kinetic nucleation theory: A new expression for the rate of homogeneous nucleation from an ideal supersaturated vapor*, J. Chem. Phys., 93 (1990), pp. 1273–1277.

[48] J. Gómez-Morales, J. Torrent-Burgués, and R. Rodríguez-Clemente, *Nucleation of calcium carbonate at different initial pH conditions*, J. Cryst. Growth, 169 (1996), pp. 331–338.

[49] J. Goodman and J. Weare, *Ensemble samplers with affine invariance*, Com. App. Math. Comp. Sci., 5 (2010), pp. 65–80.

[50] D. Goodwin, N. Malaya, H. Moffat, and R. Speth, *Cantera: An object-oriented software toolkit for chemical kinetics, thermodynamics, and transport processes*. `https://code.google.com/p/cantera`, Version 2.1a1 2013.

[51] R. Gordon, *Error bounds in equilibrium statistical mechanics*, J. Math. Phys., 9 (1968), p. 655.

[52] G. Goujon and B. Mutaftschiev, *On the crystallinity and the stoichiometry of the calcite surface*, J. Colloid Interface Sci., 57 (1976), pp. 148–161.

[53] P. Gregory, *Bayesian Logical Data Analysis for the Physical Sciences*, Cambridge University Press, New York, 2005.

[54] E. Guggenheim, *The thermodynamics of interfaces in systems of several components*, Trans. Faraday Soc., 35 (1940), pp. 397–412.

[55] G. Guisbiers, *Size-dependent materials properties toward a universal equation*, Nanoscale Res. Lett., 5 (2010), pp. 1132–1136.

[56] ——, *Schottky defects in nanoparticles*, J. Phys. Chem. C, 115 (2011), pp. 2616–2621.

[57] G. Guisbiers and G. Buchallot, *Universal size/shape-dependent law for characteristic temperatures*, Phys. Lett. A, 374 (2009), pp. 305–308.

[58] A. Gutjahr, H. Dabrinhaus, and R. Lacmann, *Studies of the growth and dissolution kinetics of the $CaCO_3$ polymorphs calcite and aragonite I. Growth and dissolution rates in water*, J. Cryst. Growth, 158 (1996), pp. 296–309.

[59] D. Hamby, *A review of techniques for parameter sensitivity analysis of environmental models*, Environ. Monit. Assess., 32 (1994), pp. 135–154.

[60] D. Harris, *A Kinetic Study of Aqueous Calcium Carbonate*, Master's thesis, Brigham Young University, Provo, U.T., 2013.

[61] D. Higdon, M. Kennedy, J. Cavendish, J. Cafeo, and R. Ryne, *Combining field data and computer simulations for calibration and prediction*, SIAM J. Sci. Comput., 26 (2004), pp. 448–466.

[62] V. Holten, D. Labetski, and M. van Dongen, *Homogeneous nucleation between water between 200 and 240 K: New wave tube data and estimation of the Tolman length*, J. Chem. Phys., 123 (2005), p. 104505.

[63] L. Holysz and E. Chibowski, *Surface free energy components of calcium carbonate and their changes due to radiofrequency electric field treatment*, J. Coll. Interf. Sci., 164 (1994), pp. 245 – 251.

[64] M. Hounslow, H. Mumtaz, A. Collier, J. Barrick, and A. Bramley, *A micro-mechanical model for the rate of aggregation during precipitation from solution*, Chem. Eng. Sci., 56 (2001), pp. 2543–2552.

[65] D. Ilievski and I. Livk, *An agglomeration efficiency model for gibbsite precipitation in a turbulently stirred vessel*, Chem. Eng. Sci., 61 (2006), pp. 2010–2022.

[66] E. Jaynes, *Probability Theory: The Logic of Science*, Cambridge University Press, New York, 2003.

[67] E. Jones, T. Oliphant, P. Peterson, et al., *SciPy: Open source scientific tools for Python*. `http://www.scipy.org`, 2001–.

[68] V. Kalikmanov, *Semiphenomenological theory of the Tolman length*, Phys. Rev. E, 55 (1997), pp. 3068–3071.

[69] D. Kashchiev, *Solution of the non-steady state problem in nucleation kinetics*, Surf. Sci., 14 (1969), pp. 209–220.

[70] ——, *Nucleation: Basic Theory with Applications*, Butterworth-Heinemann, New York, 2000.

[71] D. Kashchiev and G. van Rosmalen, *Review: Nucleation in solutions revisited*, Cryst. Res. Tech., 38 (2003), pp. 555–574.

[72] R. Kass and A. Raftery, *Bayes factors*, J. Am. Stat. Assoc., 90 (1995), pp. 773–795.

[73] K. Kelton, A. Greer, and C. Thompson, *Transient nucleation in condensed systems*, J. Chem. Phys., 79 (1983), pp. 6261–6276.

[74] M. Kennedy and A. O'Hagan, *Bayesian calibration of computer models*, J. Roy. Stat. Soc. B. Met., 63 (2001), pp. 425–464.

[75] D. Kralj, L. Brečević, and J. Kontrec, *Vaterite growth and dissolution in aqueous solution III. Kinetics of transformation*, J. Cryst. Growth, 177 (1997), pp. 248–257.

[76] D. Kralj, L. Brečević, and A. Nielsen, *Vaterite growth and dissolution in aqueous solution I. Kinetics of crystal growth*, J. Cryst. Growth, 104 (1990), pp. 793–800.

[77] ——, *Vaterite growth and dissolution in aqueous solution II. Kinetics of dissolution*, J. Cryst. Growth, 143 (1994), pp. 269–276.

[78] E. Kreyszig, *Advanced Engineering Mathematics*, John Wiley & Sons, Inc., New York, 8th ed., 1999.

[79] M. Liefvendahl and R. Stocki, *A study of algorithms for optimization of Latin hypercubes*, J. Statist. Plann. Inference, 136 (2006), pp. 3231–3247.

[80] T. Liew, J. Barrick, and M. Hounslow, *A micro-mechanical model for the rate of aggregation during precipitation from solution*, Chem. Eng. Technol., 26 (2003), pp. 282–285.

[81] I. Lifshitz and V. Slyozov, *The kinetics of precipitation from supersaturated solid solutions*, J. Phys. Chem. Solids, 19 (1961), pp. 35–50.

[82] Y.-P. Lin and P. Singer, *Effects of seed material and solution composition on calcite precipitation*, Geochim. Cosmochim. Acta, 69 (2005), pp. 4495–4509.

[83] C. Lindenberg and M. Mazzotti, *Effect of temperature on the nucleation kinetics of α L-glutamic acid*, J. Cryst. Growth, 311 (2009), pp. 1178–1184.

[84] Y. Ling and S. Mahadevan, *Quantitative model validation techniques: New insights*, Reliab. Eng. Syst. Safe., 111 (2013), pp. 217–231.

[85] Y. Ling, J. Mullins, and S. Mahadevan, *Selection of model discrepancy priors in Bayesian calibration*, J. Comp. Phys., 276 (2014), pp. 665–680.

[86] F. Liu, M. Bayarri, J. Berger, R. Paulo, and J. Sacks, *A Bayesian analysis of the thermal challenge problem*, Comput. Method Appl. Mech. Engrg., 197 (2008), pp. 2457–2466.

[87] J. S. Liu, *Monte Carlo Strategies in Scientific Computing*, Springer Series in Statistics, Springer, New York, 2008.

[88] D. J. MacKay, *Information Theory, Inference, and Learning Algorithms*, Cambridge University Press, New York, 7.0 ed., 2003.

[89] F. Manoli and E. Dalas, *Spontaneous precipitation of calcium carbonate in the presence of chondroitin sulfate*, J. Cryst. Growth, 217 (2000), pp. 416–421.

[90] D. Marchisio and R. Fox, *Solution of population balance equations using the direct quadrature method of moments*, Aerosol Sci., 36 (2005), pp. 43–73.

[91] ——, *Multiphase Reacting Flows: Modelling and Simulation*, vol. 492 of CISM International Centre for Mechanical Sciences, Springer Vienna, Vienna, 2007.

[92] D. Marchisio, R. Vigil, and R. Fox, *Implementation of the quadrature method of moments in CFD codes for aggregation–breakage problems*, Chem. Eng. Sci., 58 (2003), pp. 3337–3351.

[93] ——, *Quadrature method of moments for aggregation – breakage processes*, Simulation, 258 (2003), pp. 322–334.

[94] R. McGraw, *Description of aerosol dynamics by the quadrature method of moments*, Aerosol Sci. Tech., 27 (1997), pp. 255–265.

[95] D. Merrick, *Mathematical models of the thermal decomposition of coal*, Fuel, 62 (1983), pp. 540–546.

[96] A. Mersmann, *Calculation of interfacial tensions*, J. Cryst. Growth, 102 (1990), pp. 841–847.

[97] M. Moody and P. Attard, *Curvature-dependent surface tension of a growing droplet*, Phys. Rev. Lett., 91 (2003), p. 056104.

[98] J. Mullins, *Resource Allocation for Uncertainty Quantification and Reduction*, PhD thesis, Vanderbilt University, Nashville, Tennessee, December 2014.

[99] K. Neoh and R. Gannon, *Coal volatile yield and element partitioning in rapid pyrolysis*, Fuel, 63 (1984), pp. 1347–1352.

[100] A. Nielsen, *Kinetics of Precipitation*, Pergamon Press Ltd., New York, 1964.

[101] S. Niksa and A. Kerstein, *FLASHCHAIN theory for rapid coal devolatilization kinetics. 1. Formulation*, Energ. Fuel, 5 (1991), pp. 647–665.

[102] C. Noguera, B. Fritz, A. Clément, and A. Baronnet, *Nucleation, growth, and aging scenarios in closed systems I: A unified mathematical framework for precipitation, condensation, and crystallization*, J. Cryst. Growth, 297 (2006), pp. 180–186.

[103] ——, *Nucleation, growth, and aging scenarios in closed systems II: Dynamics of a new phase formation*, J. Cryst. Growth, 297 (2006), pp. 187–198.

[104] W. Oberkampf and C. Roy, *Verification and Validation in Scientific Computing*, Cambridge University Press, New York, 2010.

[105] W. Oberkampf and T. Trucano, *Verification and validation in computational fluid dynamics*. SAND2002 - 0529, March 2002.

[106] T. Ogino, T. Suzuki, and K. Sawada, *The formation and transformation mechanism of calcium carbonate in water*, Geochim. Cosmochim. Acta, 51 (1987), pp. 2757–2767.

[107] T. Oliver, G. Terejanu, C. Simmons, and R. Moser, *Validating predictions of unobserved quantities*, Comput. Method Appl. Mech. Engrg., 283 (2015), pp. 1310–1335.

[108] W. Ostwald, *Studien über die bildung und umwandlung fester körper*, Z. Phys. Chem., 22 (1897), pp. 289–330.

[109] L. Plummer and E. Busenberg, *The solubilities of calcite, aragonite, and vaterite in CO_2-H_2O solutions between 0 and 90°C, and an evaluation of the aqueous model for the system $CaCO_3$-CO_2-H_2O*, Geochim. Cosmochim. Acta, 46 (1982), pp. 1011–1040.

[110] L. Plummer, T. Wigley, and D. Parkhurst, *The kinetics of calcite dissolution in CO_2-water systems at 5° to 60°C and 0.0 to 1.0 atm CO_2*, Am. J. Sci., 278 (1978), pp. 179–216.

[111] O. Pokrovsky, S. Golubev, J. Schott, and A. Castillo, *Calcite, dolomite and magnesite dissolution kinetics in aqueous solutions at acid to circumneutral pH, 25 to 150°C and 1 to 55 atm pCO_2: New constraints on CO_2 sequestration in sedimentary basins*, Chem. Geol., 265 (2009), pp. 20–32.

[112] A. Radha, T. Forbes, C. Killian, P. Gilbert, and A. Navrotsky, *Transformation and crystallization energetics of synthetic and biogenic amorphous calcium carbonate*, Proc. Natl. Acad. Sci. U.S.A., 107 (2010), pp. 16438–16443.

[113] D. Ramkrishna, *Population Balances: Theory and Applications to Particulate Systems in Engineering*, Academic Press, New York, 2000.

[114] A. Randolph and M. Larson, *Theory of Particulate Processes: Analysis and Techniques of Continuous Crystallization*, Academic Press, Inc., New York, 2nd ed., 1971.

[115] C. Rasmussen and C. Williams, *Gaussian Processes for Machine Learning*, The MIT Press, Cambridge, Massachusetts, 2006.

[116] A. Raue, C. Kreutz, T. Maiwald, J. Bachmann, M. Schilling, and Klingmüller, *Structural and practical identifiability analysis of partially observed dynamical models by exploiting the profile likelihood*, Bioinformatics, 25 (2009), pp. 1923–1929.

[117] J. Riihimaki and A. Vehtari, *Gaussian processes with monotonicity information*, 13th AISTATS, 9 (2010).

[118] J. Rodriguez-Blanco, S. Shaw, and L. Benning, *The kinetics and mechanisms of amorphous calcium carbonate (ACC) crystallization to calcite, via vaterite*, Nanoscale, 3 (2011), pp. 265–271.

[119] E. Ruiz-Agundo, C. Putnis, C. Rodriguez-Navarro, and A. Putnis, *Effect of pH on calcite growth at constant $a_{ca^{2+}}/a_{co_3{}^{2-}}$ ratio and supersaturation*, Geochim. Cosmochim. Acta, 75 (2011), pp. 284–296.

[120] T. Russi, *Uncertainty Quantification with Experimental Data and Complex System Models*, PhD thesis, University of California, Berkeley, 2010.

[121] T. Russi, A. Packard, R. Feeley, and M. Frenklach, *Sensitivity analysis of uncertainty in model prediction*, J. Phys. Chem. A, 112 (2008), pp. 2579–2588.

[122] T. Russi, A. Packard, and M. Frenklach, *Uncertainty quantification: Making predictions of complex reaction systems reliable*, Chem. Phys. Lett., 499 (2010), pp. 1–8.

[123] J. Sacks and W. Welch, *Gaussian Process model - Tutorial 2.* SAMSI Surrogate Models Working Group based upon 2010 short course, 2010.

[124] J. Sacks, W. Welch, T. Mitchell, and H. Wynn, *Design and analysis of computer experiments*, Stat. Sci., 4 (1989), pp. 409–423.

[125] M. Sander, R. Paterson, R. Braumann, A. Raj, and M. Kraft, *Developing the PAH-PP soot particle model using process informatics and uncertainty propagation*, P. Combust. Inst., 33 (2011), pp. 675–683.

[126] S. Schlesinger, *Terminology for model credibility*, Simulation, 32 (1979).

[127] B. Schroeder, D. Harris, S. Smith, and D. Lignell, *Theoretical framework for multiple-polymorph particle precipitation in highly supersaturated systems*, Cryst. Growth Des., 14 (2014), pp. 1756–1770.

[128] D. Sivia, *Data Analysis: A Bayesian Tutorial*, Oxford University Press, New York, 1996.

[129] D. Smoot and P. Smith, *Coal Combustion and Gasification*, The Plenum Chemical Engineering Series, Plenum Press, New York, 1985.

[130] O. Söhnel and J. Mullin, *Precipitation of calcium carbonate*, J. Cryst. Growth, 60 (1982), pp. 239–250.

[131] P. Solomon, D. Hamblen, R. Carangelo, M. Serio, and G. Deshpande, *General model of coal devolatilization*, Energ. Fuel, 2 (1988), pp. 405–422.

[132] A. Stack and M. Grantham, *Growth rate of calcite steps as a function of aqueous calcium-to-carbonate ratio: Independent attachment and detachment of calcium and carbonate ions*, Cryst. Growth Des., 10 (2010), pp. 1409–1413.

[133] C. Steefel and P. van Cappellen, *A new kinetic approach to modeling water-rock interaction: The role of nucleation, precursors, and ostwald ripening*, Geochim. Cosmochim. Acta, 54 (1990), pp. 2657–2677.

[134] W. Stumm and J. Morgan, *Aquatic Chemistry: Chemical Equilibria and Rates in Natural Waters*, Environmental Science and Technology, John Wiley & Sons, Inc., New York, 3rd ed., 1996.

[135] L. Swiler, B. Adams, and M. Eldred, *Model calibration under uncertainty: Matching distribution information*, SAND Report 2008-0632A, Sandia National Laboratories, Albuquerque, NM 87185, 2008.

[136] B. Thacker, *ASME Standards Committee on Verification and Validation in Computational Solid Mechanics*, Report, The American Society of Mechanical Engineers Council on Codes and Standards, September 2001.

[137] R. C. Tolman, *The effect of droplet size on surface tension*, J. Chem. Phys., 17 (1948), pp. 333–337.

[138] M. Tröster, M. Oettel, B. Block, P. Virnau, and K. Binder, *Numerical approaches to determine the interfacial tension of curved interfaces from free energy calculations*, J. Chem. Phys., 136 (2012), p. 064709.

[139] University of Utah Institute for Clean and Secure Energy, *ARCHES*. `www.uintah.utah.edu`.

[140] M. van der Hoef, M. van Sint Annaland, I. A. Andrews, S. Sundaresan, and J. Kuipers, *Multiscale modeling of gas-fluidized beds*, Adv. Chem. Eng., 31 (2006), pp. 65–149.

[141] A. van Giessen and E. Blokhuis, *Direct determination of the Tolman length from bulk pressues of liquid drops via molecular dynamics simulations*, J. Chem. Phys., 131 (2009), p. 164705.

[142] C. van Oss, *Interfacial Forces in Aqueous Media*, Taylor & Francis Ltd, Boca Raton, FL, 2nd ed., 2006.

[143] L. Wang and R. Fox, *Comparison of micromixing models for CFD simulation of nanoparticle formation*, AIChE J., 50 (2004), pp. 2217–2232.

[144] B. Williams, D. Higdon, J. Gattiker, L. Moore, M. McKay, and S. Keller-McNulty, *Combining experimental data and computer simulations, with an application to flyer plate experiments*, Bayesian Anal., 1 (2006), pp. 765–792.

[145] M. Wolthers, G. Nehrke, J. Gustafsson, and P. van Cappellen, *Calcite growth kinetics: Modeling the effect of solution stoichiometry*, Geochim. Cosmochim. Acta, 77 (2012), pp. 121–134.

[146] W. Wu and G. H. Nancollas, *Determination of interfacial tension from crystallization and dissolution data: A comparison with other methods.*, Adv. Colloid Interfac., 79 (1999), pp. 229–79.

[147] X. You, T. Russi, A. Packard, and M. Frenklach, *Optimization of combustion kinetic models on a feasible set*, P. Combust. Inst., 33 (2011), pp. 509–516.

[148] C. Yuan, F. Laurent, and R. Fox, *An extended quadrature method of moments for population balance equations*, J. Aerosol Sci., 51 (2012), pp. 1–23.

[149] J. Zeldovich, *Theory of the formation of a new phase cavitation*, Zh. Eksp. Teor. Fiz., 12 (1942).