UNCERTAINTY QUANTIFICATION AND DECISION MAKING IN

HIERARCHICAL DEVELOPMENT OF COMPUTATIONAL MODELS

By

Angel Urbina

Dissertation

Submitted to the Faculty of the

Graduate School of Vanderbilt University

in partial fulfillment of the requirements

for the degree of

DOCTOR OF PHILOSOPHY

in

Civil Engineering

August, 2009

Nashville, Tennessee

Approved:

Professor Sankaran Mahadevan

Professor Prodyot K. Basu

Professor Gautam Biswas

Professor Bruce Cooil

Professor Thomas L. Paez

To my wonderful wife, Ruth and my beautiful daughters, and to my
teachers for their patience and guidance

ACKNOWLEDGMENTS

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER I

INTRODUCTION

1.1 Motivation

The reliability of high consequence systems, such as weapon systems, has been traditionally established by testing individual systems and verifying that their performance is within some acceptable limits. Although full scale testing is currently not feasible for some full systems under actual use environments, some limited testing is often available for components, assemblies (i.e. groups of components) and a very limited number of tests of the full system in other use environments or in laboratory controlled tests. Modeling and simulation fill the gap left by the lack of full scale testing for the actual use environments. Because component level data are cheaper and easier to obtain relative to the system data, it is advantageous to have the ability to build individual models of the components and/or assemblies using the available data and incorporate them into the system level model. This leads to a hierarchical approach to building system level models and consequently the uncertainty in the system level model is a function of the component level data and of the knowledge not captured in the component level data. Furthermore, because tests cannot be performed for many actual use environments, the model is required to extrapolate beyond the data it was developed from. To establish confidence in an extrapolated model prediction, sources of uncertainty must be identified, quantified and propagated through the model to the response quantity of interest at the system level.

One formal framework to establish confidence in an extrapolated system level model response within the context of nuclear weapons was proposed at Sandia National Laboratories (referred to as Sandia hereafter) and it is referred to as Risk Informed Decision Analysis (RIDA). A central statement of RIDA which provides the main motivation for this research states (Pilch et al., 2006):

> "Whatever mathematical form an application of Risk Informed Decision Analysis (RIDA) to a stockpile lifecycle decision might take, it requires that **all uncertainties** be identified and characterized. This includes the separate **quantification** of both **variability** (i.e., aleatoric uncertainty) and **lack-of-knowledge uncertainty** (i.e., epistemic uncertainty), as well as definitions of "other factors" and quantified characterizations of their individual contributions to uncertainty. RIDA also requires attention to uncertainties in requirements and decision criteria, such as definitions of performance thresholds that are fundamental to the decision making. In addition, RIDA requires complete transparency of all the information to make the decision process understandable, traceable, and reproducible (documented)."

For the purpose of this dissertation, it will be assumed that the statement "… all uncertainties be identified and characterized" refers to all relevant uncertainties and not the whole universe of what is uncertain about a system. In this context, aleatoric uncertainty or variability is irreducible and in most cases can be characterized via a probability density function. These arise due to inherent variations in materials, assembly of components and test conditions. Epistemic uncertainty is due to lack of knowledge about a particular behavior in a system. These arise due to non-existent or incomplete data regarding a material behavior or a certain phenomenon, uncertainty in the choice of a computational model, the coupling of two distinct phenomena for which no information is available, and the probability density form and statistics chosen to represent a random variable. Epistemic uncertainty is reducible in the sense that with more information one

could gain increased insight into the behavior of a particular variable. If we seek to make our models a closer representation of an actual system, these types of uncertainties must be included in the analysis. Returning to the RIDA statement, notice that it allows for any mathematical framework to be used, whether classical or Bayesian probability methods, evidence theory or others, provided that it can be documented and defensible. This statement also requires the identification, quantification and propagation of uncertainty throughout the process. Dealing with the different types of uncertainties present in the system is a central topic of this research.

A central concept in this process of RIDA is the Quantification of Margins and Uncertainty (QMU). In Pilch et al, 2006, QMU is defined as:

> "QMU is a decision-support methodology for complex technical decisions centering on performance thresholds and associated margins for engineered systems that are made under conditions of uncertainty. RIDA does not base its decision outcomes solely on the results of QMU. Rather, QMU provides only part of the input into the decision process."

QMU is thus the methodology that requires quantification of performance thresholds and margins, as well as the associated uncertainty in their evaluation. Pilch et al, 2006 does not provide a mathematical formulation of how this quantification should be done but it only spells out the steps or elements that such quantification should have. This provides an opportunity to use tools from well-established fields, such as reliability analysis, to address this.

At Sandia, there has been an emphasis on developing models of components from first principles, calibrating them from simple exploratory experiments, validating them relative to a different set of experiments and then using them within a more complex

model. What was described above is a hierarchical approach to building a system level model and has been repeatedly used at Sandia and other Department of Energy research laboratories. It basically is a construction of a complex system model by using a building block approach that incorporates simpler component based models and couples them together. This research proposes to use any available data to augment the knowledge base used to infer the uncertainty in the parameters of a model. By using Bayesian updating techniques and Bayes networks, it should be possible to incorporate the available data at multiple levels, update the model parameters and make model predictions to reflect the new information that was previously not available to the other individual levels. This methodology has been explored initially in Rebba (2005) for a system consisting of two levels of complexity and the current research extends this work to a 2-component system.

In this research, the aim is to quantify and propagate uncertainty and evaluate confidence in the extrapolated response of the model prediction at actual use conditions. To demonstrate the methodologies developed in this research, an example problem developed at Sandia is used and it is described in Chapter 3. The proposed problem has the following characteristics:

1. It is a 2-component problem where one branch involves a mechanical joint and the other, an encapsulating foam. Both are energy dissipating mechanisms.

2. It is a multi-level problem where the phenomenon observed at the lowest level is assumed to be present at subsequent levels, i.e. damping in the joints and foam is assumed to be present at all levels. The degree to which the damping at one level is similar to the damping at another is a source of uncertainty and thus, the

relationship of energy dissipation at one level to energy dissipation at another level may be difficult to establish.

3. The individual component branches, joint and foam, converge to a system level hardware where the two couple together. The interaction of the two physics is a potential source of epistemic uncertainty because this had not been previously tested.

4. Experimental data consists of repeated tests on several, nominally identical hardware systems. These are intended to quantify the variability inherent in a physical system.

5. Finite element models for all levels are built and verified (i.e. mesh convergence, time step convergence and PDE solver tolerance are calibrated to generate stable answers) and used to simulate a particular behavior of the physical hardware they represent. The model parameters have been calibrated from simple, discovery experiments aimed at isolating the particular physical phenomena that the model is meant to represent.

1.2 Research Objectives

The main objective of this work is to incorporate various sources of uncertainty into an analysis framework that combines information, in a probabilistic manner, to quantify and propagate uncertainty to a system level response. Information is contained at different levels of complexity and it also comes from different physical sources (i.e. different components). The work presented here establishes a framework that allows

various sources of information to be combined in a probabilistic manner and incorporates different sources of uncertainty. To address the main focus of this research, the following objectives are proposed:

1. Quantification and propagation of aleatoric uncertainty (variability) in a hierarchically built system model;

2. Quantification and propagation of both aleatoric and epistemic uncertainty (lack of or incomplete knowledge) in a hierarchically built system model;

3. Quantification of margins and uncertainties (QMU) in a system level response to support a risk informed decision analysis; and

4. Resource allocation using QMU.

Research objective 1 establishes the baseline framework to propagate and quantify sources of uncertainty. Uncertainty quantification analysis will be done on a hierarchical system model which incorporates sources of aleatoric uncertainty within the context of a multi-level, multi-component problem using a Bayes network as the main analysis framework. The Bayes network is chosen because it conveniently allows the modeling of both causal-type relationships and statistical dependencies between sets of data within a probability framework based on conditional probabilities between the various nodes of the network. This framework also allows any available experimental data to be incorporated into the analysis and updates the conditional probabilities. Once updated, the posterior probabilities of all nodes in the network can be obtained. The extrapolated response of the full system is also investigated. This objective extends the work of Rebba (2005) by considering a multi-level, multi-component system and

quantifies uncertainty in the extrapolated quantity. A key issue to address is the actual implementation of this methodology to perform fast and efficient computation.

Research objective 2 extends the methodology described in objective 1 to incorporate sources of epistemic uncertainty into the Bayes network framework developed in objective 1. For this objective, the main difference is that some of the input parameters to a model would be specified in terms of bounds on a parameter given by experts and not by full probability density functions. An approach that is currently under development and described in McDonald et al, (2009a) and in Venkataraman and Wilson (1987) will be implemented in this research. This approach starts by calculating the sample mean and the next three central moments of intervals representing the epistemic uncertainty as described in Ferson et al (2007) and implemented in McDonald et al (2009 b,c). These moments are then used to estimate the parameters of a Johnson family of distributions by any one of four methods presented in DeBrota et al (1998). The Johnson distributions will be incorporated in the Bayes network in a manner similar to objective 1.

Research objective 3 implements the quantification of margins and uncertainty (QMU) approach using the uncertainty in the extrapolated model response obtained in objectives 1 and 2. In this research an approach similar to reliability analysis will be used to implement QMU. It is based on calculating the amount of overlap between a system level performance and a given requirement. This overlap defines the probability of failure of the system and it is used as a comparison metric for the purpose of this dissertation. This approach uses the uncertainty quantification results obtained in Objectives 1 and 2 and a given performance requirement to estimate the probability of failure. A comparison of the results when aleatoric and epistemic uncertainty is included in the analysis versus

only aleatoric uncertainty is presented. Additionally, a comparison of the failure metrics obtained with different methodologies to model epistemic uncertainty is presented.

Research objective 4 ties together the previous three objectives and considers the resource allocation in terms of improving the confidence in the system model. Based on the required confidence level and the calculated confidence based on the model response, an assessment will be made regarding whether or not the system is "certified" relative to the requirements. If the answer is no, then based on sensitivity analysis of the contribution of various nodes of the Bayes network to the overall model confidence, resource allocation guidance to achieve the required level of confidence could be given. Examples of resources are more experiments at a certain level, more model simulations, model refinement, reduction in uncertainty (particularly, epistemic uncertainty) in the model parameters, and reduction in modeling errors due to solution approximation, model form error and others. Different actions have their associated costs. Various scenarios to achieve the required confidence are examined in this study.

These objectives treat the various sources of uncertainty in a quantitative way, and when available, existing methods for addressing each objective are implemented and extended to multi-level models.

1.3 Organization of Dissertation

This dissertation is organized as follows.

Chapter 2 provides background material and terminology relevant to this work. The formal weapon assessment process is briefly explained as well as the relevance of this

work to the current assessment process. The various technological challenges are noted in this chapter as well.

Chapter 3 presents the specific example problem to be solved. This is a Sandia National Laboratories relevant problem and details of both the testing and the modeling aspects are presented here. Determination of the relevant quantities of interest at each level of complexity and for each of the components in the problem as well as constructing surrogate models as an enabling technology are discussed within this chapter.

Chapter 4 considers the baseline case where sources of uncertainty in the model input are considered aleatoric and treated in a probabilistic framework. This chapter introduces the basics of Bayesian analysis which includes Bayes networks and implements the proposed methodology to the problem described in Chapter 3.

Chapter 5 extends the methodology developed and implemented in Chapter 4 to cases were both epistemic and aleatoric uncertainties are present. The implementation of various methodologies to accommodate the probabilistic treatment of epistemic uncertainty is described in this chapter and a comparison of various sets of results for the system level quantity of interest are presented here.

Chapter 6 implements a reliability based approach to addressing QMU. This will be done by calculating the area of overlap between a system level response and a given threshold. This area is traditionally refer to as the probability of failure of a system and it is proposed in this work to be used as a metric of performance at the system level which incorporates sources of uncertainty in both the system level response and in the threshold.

Chapter 7 presents a methodology for resource allocation toward improvement of model confidence which incorporates analyses shown in Chapter 5 and Chapter 6. An optimization based solution is proposed to reduce the uncertainty in the decision metric of interest. A potential use of the results shown in this chapter will be to allow decision-makers a way to allocate resources (computational, testing, monetary, etc) to increase the confidence in the system level prediction.

Finally, the Chapter 8 presents general conclusions as well as recommendations for future work.

CHAPTER II

BACKGROUND AND RELEVANT LITERATURE

Four main research objectives were identified in Chapter 1 to address the overall goal of uncertainty quantification in a hierarchically built system model. The objectives are (1) Quantification and propagation of aleatoric uncertainty (variability), (2) Quantification and propagation of both aleatoric and epistemic uncertainty (lack of or incomplete knowledge), (3) Quantification of margins and uncertainties in system level prediction (QMU) to support a risk informed decision analysis and (4) Decision making for resource allocation towards improvement of model confidence and system performance. A review of the state-of-the-art for the various supporting methodologies used in this research is presented in the following sections.

The sections of this chapter are organized as follows. Section 2.1 provides the background of the Risk Informed Decision Analysis (RIDA) paradigm which is the central motivating concept for this dissertation. Section 2.2 defines the concept of quantification of margins and uncertainties and its application to a modeling and simulation approach to quantifying uncertainty and establishing confidence in an extrapolated model prediction of a system. Section 2.3 considers some pros and cons of applying Bayesian analysis to address uncertainty quantification problems. This research proposes the use of Bayesian analysis as a tool to integrate observed data and prior knowledge about a quantity of interest to infer the distribution of this quantity given the observed data (i.e. the posterior distribution estimate of a quantity of interest). In doing

so, uncertainty is propagated through the hierarchical model to the system level. It is this uncertainty that is then used in the QMU analysis to estimate the confidence in the system level prediction. As stated in Chapter 1, QMU is the methodology that supports RIDA. Section 2.4 examines the sources of uncertainty and their treatment, particularly the probabilistic treatment of epistemic uncertainty.

2.1 Risk Informed Decision Analysis (RIDA)

This section starts with the statement of what qualities RIDA should posses according to Pilch et al. (2006) which was mentioned in Chapter 1. This statement, to say the least, is a complex set of requirements that incorporates abstractions that may or may not be attainable. Nevertheless, it is a challenge that is presented to the research community and should be addressed in a systematic way. The salient features of this statement are the text underlined and in bold typeface. The first point is the identification of all uncertainties present in a given system. This could be a monumental task, particularly for a complex system, because by definition, uncertainty is all that is not known for a fact. This implies that the categories of things that could be unknown, becomes known. This statement will be interpreted to mean that all sources of **relevant** uncertainty that substantially contribute to the uncertainty of system behavior will be identified which is more conducive to a real-life application. This statement immediately leads to the second highlighted point which is the quantification of both aleatoric and epistemic uncertainty. This topic will be extensively covered in section 2.3. The final issue is the attention to uncertainties in requirements and decision criteria which

ultimately leads to a final decision regarding the usability of the high consequence systems. This process is known as certification.

The developments in this work relate to UQ, QMU, RIDA and other modeling and assessment activities at DOE laboratories. Formally, the current annual certification process is a series of formalized reviews, conducted each year with multiple participants from various government and contractor organizations, culminating in a written certification letter from the US Secretaries of Defense and Energy to the President of the United States that the stockpile is safe and reliable in the absence of underground testing. The process serves to provide the President of the United States, and also the Secretary of Defense through Department of Defense participation, a measure of confidence that the nuclear deterrent is still safe and militarily effective (Perkins, 2000). The certification process is shown schematically in Figure 2.1. (Aloise, 2007).
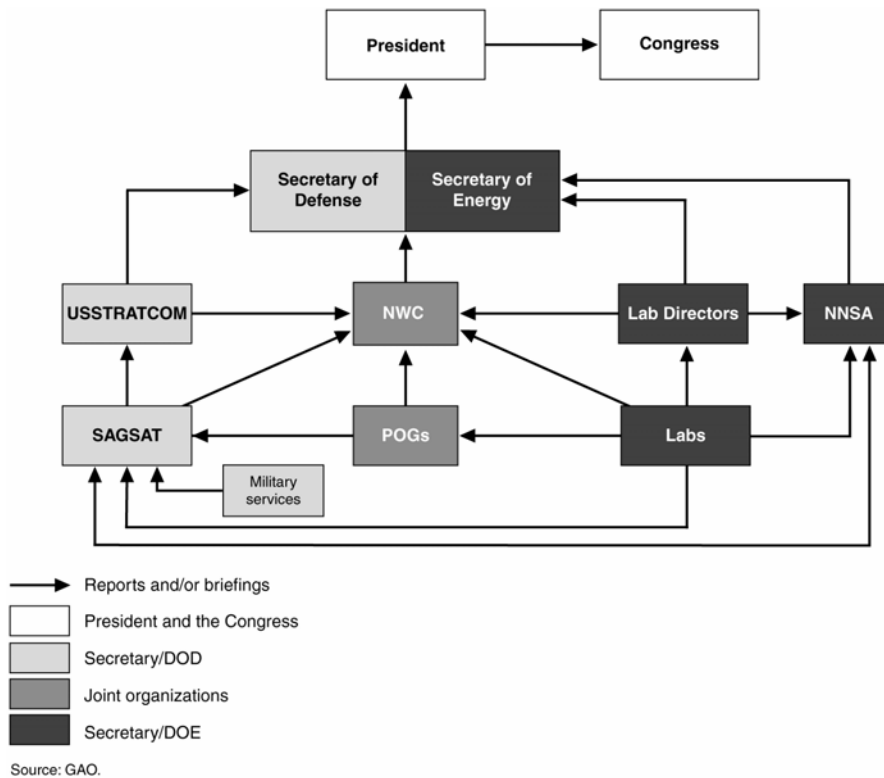
Figure 2.1. Annual Certification Process of Nuclear Weapons Stockpile

This certification process is what RIDA supports and the work at Sandia is integral to the process because it provides the analysis of the components and this eventually leads to a decision by other members in the certification hierarchy. The purpose of this certification is to establish a confidence in the stockpile which has a quantitative and a qualitative aspect. This is noted in the 1999 Foster-panel's report to Congress (Foster, 1999). It describes "'confidence in the stockpile' as having two dimensions. The first, being quantitative, derives from scientific assessments and surveillance of the systems. It is based largely on a sufficient degree of thoroughness and probabilities calculated as a result of a wide array of tests. The second dimension of confidence centers on judgment; it is based on trust in the ability of the people, methods, and tools available to find,

14

assess, and fix potential problems in the stockpile. Because of its qualitative nature and heavy reliance on judgment, it is this second facet which can be most perplexing. In fact, as we look closer at the overall process, we will find that once all the caveats and assumptions are accounted for in "quantitative testing," here too, a great deal of faith is placed in judgment." Furthermore, under the current ban on testing, the first dimension mentioned in this report is now replaced by modeling and simulation which brings additional sources of incertitude to the problem. It is this issue which is addressed in the present research.

2.2 Quantification of Margins and Uncertainties (QMU)

A central concept in this process of risk informed decision analysis (RIDA) is the Quantification of Margins and Uncertainty (QMU) (Eardley, 2005; Sharp and Wood-Schultz, 2003). QMU is defined by Sharp and Wood-Schultz as "a framework that captures what we do and do not know about the performance of a nuclear weapon in a way that can be used to address risk and risk mitigation." Pilch et al (2006) define it as "a decision-support methodology for complex technical decisions centering on performance thresholds and associated margins for engineered systems that are made under conditions of uncertainty". Pilch et al (2006) further state: "RIDA does not base its decision outcomes solely on the results of QMU. Rather, QMU provides only part of the input into the decision process". QMU is thus the mathematical methodology that quantifies these thresholds and margins, as well as the associated uncertainty in their evaluation. QMU is applied in a decision-making context, addressing the ability to meet design, qualification, or life-cycle performance requirements. Because of the programmatic constraints of cost

15

and schedule, there are often significant uncertainties due to lack of knowledge associated with the use of a model. QMU, particularly when using models whose uncertainties are dominated by lack of knowledge issues, has the technical dimensions of quantitative risk assessment (QRA). From the perspective of QRA, risk can be defined in terms of the Kaplan and Garrick (1981) risk triplet:

1. *Scenario identification* – What can happen?

2. *Likelihood of scenarios* – How likely is it to happen?

3. *Consequence of scenarios* – What are the consequences if it does happen?

A fourth component has always been an important factor in the use of QRA and will be an important factor in the application of QMU to the stockpile:

4. *Credibility* – How much confidence do you have in the answers to the first three questions?

The guidance for our QMU framework must be formulated in a manner that can easily address these four questions. Performance or safety requirements establish the metrics by which "consequence" can be measured in the context of a particular application. These metrics can be specified in a deterministic way (for example a given threshold (temperature) at which a component fails) or in a probabilistic fashion, such as a limit on the probability of a threshold event being surpassed. An illustration of margin and uncertainty in both performance requirement and predicted behavior of a system, and the margin between them is shown in Figure 2.2.
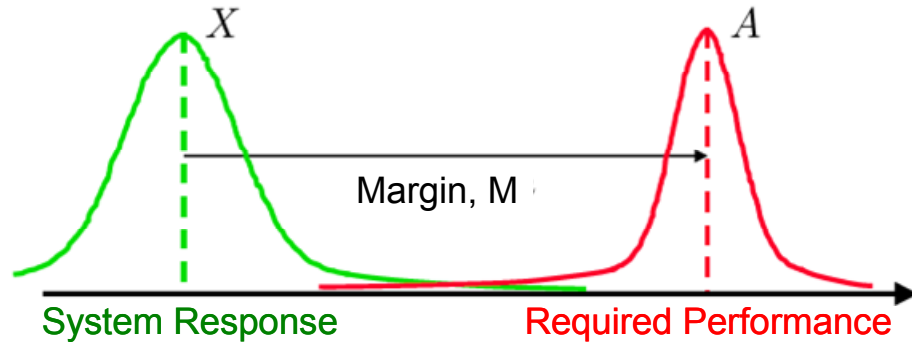
Figure 2.2. Uncertainty in predicted system response (X) and required performance (A) and the margin between them

The case shown in Figure 2.2 is the most general case which includes both uncertainty (represented by a probability density function (PDF)) in the predicted system response, X, and in the performance or requirement measure, A. One factor that complicates the application of this QMU formulation to the stockpile is the fact that the system response is an extrapolated response to an environment that cannot be tested; therefore, the system response has only the confidence that is built through the hierarchical modeling of the full system. This implies that the amount of uncertainty represented by the system PDF might not be perfectly representative of the full system when subjected to the actual use environment. In addition, the system response measure PDF approximation might not be unique, especially in the presence of epistemic uncertainty. This PDF could be replaced by either a family of PDFs or some bounding PDFs that reflect the lack of knowledge present in the system. This complicates the analysis further by making the definition of the margin a random variable as well.

2.3 Bayesian Approach to Uncertainty Quantification

One of the first issues to address in this research is the reason why we should use a Bayesian approach to solve the given Sandia problem. This research proposes the use of Bayesian methods, particularly a Bayes network as a tool to integrate observed data and prior knowledge, and quantify the uncertainty in system-level response prediction. The uncertainty in the system level response is used in the QMU analysis to estimate the confidence in the system level prediction. As stated previously, QMU is the methodology that supports RIDA and thus the connection between the proposed Bayesian analysis to the motivating statement of RIDA is through the quantification of output uncertainty.

It is noted that a team at Sandia (which included the author) used the same demonstration problem that is used here and implemented a classical probabilistic approach to model the hierarchical system (Urbina et al., 2006). It is important to note that this research does not emphasize using one approach over another nor does it attempt to compare two methodologies to determine which one is superior. The approach used in the Sandia study was based on estimating parameters of a model and their corresponding uncertainty at the most fundamental level and use them to calculate the response of higher complexity models. Uncertainty was propagated to the system level model via the finite element models that mapped the model parameters to a given behavior and using Monte Carlo sampling. Although the ultimate objective in that study was not decision making but model validation, it still required uncertainty quantification which is a part of QMU. Deficiencies found in the model were addressed by modifying the finite element model to account for discrepancies between the actual data and the model predictions but

only by using very limited amount of information. From this stand point, the current research includes all the available data to update the uncertainty in the various component models and in the system level model. One could argue that if the data is available at any level, one should be able to use it to create the best predictive models possible. This is a major difference in the approach taken in this research relative to the Sandia approach where only the data at the lowest level was used to quantify uncertainty in the parameters and no further updating of the parameters was made using higher level information. With this comparison in mind, focus is now given to why Bayesian methods are a viable alternative.

The essentials of Bayesian thinking are summarized by Gill (2002):

1. Specify a probability model that includes some prior knowledge about the parameters if available for unknown parameter values

2. Update knowledge about the unknown parameters by a conditional probability of this parameter given the observed data

3. Evaluate the fit of the model to the data and the sensitivity of the conclusions to the assumptions.

A recent text on Bayesian analysis (Gelman et al., 2004) concludes that the principles of Bayesian inference are now well established and the only issue left to address is the efficient implementation of Bayesian inference to real-world problems. This speaks to the maturity of the Bayesian methods and now the need to fully incorporate them into the application domain. Published reports such as Robinson (2001) and Williams et al. (2006) demonstrate the potential for using Bayesian methods in high consequence

environments such as at Sandia National Laboratories and Los Alamos National Laboratories.

Markov Chain Monte Carlo (MCMC) simulation has allowed Bayesian analysis to become tractable and a viable approach to solving engineering problems. Bayesian models are cursed by inference problems that were analytically intractable due to the high-dimension integral calculations that need to be solved. This was addressed with the pioneering work on MCMC of Metropolis et al. (1953), Hastings (1970), Peskun (1973) and Gelfand and Smith (1990) and has culminated with the implementation of a general purpose MCMC software, WinBUGS (Spiegelhalter, 2003). In a nutshell, the basic principle behind MCMC is that if an iterative chain of consecutive random variable realizations can be set up and run long enough, then empirical estimations of the integral quantities of interest can be obtained for the later chain values (Gill, 2002). Two potential issues arise with the use of MCMC techniques. One is the convergence of the chain itself and a related issue is the number of iterations needed for convergence. Both Gibbs and Metropolis-Hasting sampling algorithms guarantee convergence (Gill, 2002). However, this still leaves the issue of number of iterations needed for convergence which is directly related to computational expense and addresses the second reason why Bayesian methods have become a feasible option.

When using Bayesian methods that use MCMC to establish a posterior distribution of a quantity of interest, there is always the need to evaluate a function that relates some input to the quantity of interest. In applications such as those performed at Sandia, these functions are formulated in terms of complex and large finite element models that run on the order of days or weeks per function evaluation. If there is a need

to run thousands of simulations in order to establish convergence of the chain, any potential benefit gained by MCMC would be completely outweighed by this large computation expense. Surrogate models are powerful tools to address this potential pitfall. Surrogate models can be simple linear approximations, artificial neural networks, splines or Gaussian process models (a.k.a Kriging models, see McFarland, 2008).

In McFarland (2008), the use of Bayesian updating to calibrate the parameters of a model in a multi-level problem was investigated. This research used data from the highest level of complexity to calibrate a model's parameters that described the controlling physics of the problem. To expedite the updating of the model parameters, Gaussian process models were used in lieu of the full finite element model. Although some runs of the finite element model were still necessary to train the Gaussian process model, they were on the order of 10-20 versus the thousands needed for Bayesian updating, thus offering a major computational efficiency. One major difference between McFarland (2008) and the current research is in the way the model parameters are updated. In the current work, data from all levels of complexity, not just from the highest level are used. This incorporates all available information which in the particular case of nuclear weapons is more plentiful at the component level then at the full system level. Although McFarland (2008) was successful in predicting the behavior of the system in an extrapolated sense, it left the question unanswered as to what a decision maker could do if little (i.e. one) or no information is available at the system level and yet a decision needs to be made. The current research addresses this question.

2.4 Types of Uncertainties and their Treatment

Before any ideas can be put forth to address the propagation of uncertainties in a hierarchical model when both aleatoric and epistemic uncertainties are present, it is worthwhile to define each type, see how they come about and how they have been treated.

Aleatoric Uncertainty

This type of uncertainty is also referred to as variability, irreducible uncertainty, inherent uncertainty or stochastic uncertainty. This is the type most commonly associated with variability due to hardware-to-hardware and experimental setup-to-setup variability of nominally identical systems. It is also associated with material properties data and loading data. In general, a statistically significant database, fully relevant to the application is available. For the most part, a probabilistic interpretation can be assigned to input and output variables and all the machinery associated with probability theory can be used to propagate and analyze this type of uncertainty. Techniques for quantification and propagation of this type of uncertainty have been well established for many years and therefore, will not be further examined in this research.

Epistemic Uncertainty

This type of uncertainty is also known as reducible uncertainty, subjective uncertainty or lack of knowledge. Some areas where this arises are when alternate plausible models are available, cases where there is non-existent, sparse, incomplete, or inconsistent experimental data, model approximations, expert elicitation that expresses

subjective rather than data based on observations and, in general, where there is lack of information about the behavior of a system. It is found in the literature (Oberkampf, 2000; Sentz and Ferson, 2002; RESS, 2004) that this type of uncertainty is treated by two types of methods:

Non-Probabilistic Methods

- Evidence (Dempster-Shafer) theory

- Possibility theory

- Fuzzy set theory

- Interval analysis

Probabilistic Methods

- Bayesian approach or classical probability approach using transformations of bounds to probability density functions

It is important to note that both types of approaches have their pros and cons and this research will not make an attempt to demonstrate why one approach is better then any other but simply use an approach that is more suitable to be implemented within the overall analysis framework.

To begin this review, the proceedings from Sandia National Laboratories' workshop on alternative representation of epistemic uncertainty (RESS, 2004) provide a comprehensive exposition of techniques to treat problems with two types of variables, one type described by probability distributions and the other type described by interval data. This forum presented various approaches, both probabilistic and non-probabilistic.

Historically, probability theory has been used to represent epistemic uncertainty (see Apostolakis, 1990 and Parry; Winter, 1981 and Wu et al., 1990) from which a separation of aleatoric and epistemic uncertainty involves two probability spaces, one for each type of uncertainty. However, many have expressed concern about modeling epistemic uncertainty via probability density functions with the main issue being the implication of a higher resolution of knowledge than what is really present (Helton, 2006). Evidence theory has been proposed as an alternative to probability theory since it doesn't make assumptions regarding the distribution of the variables described by intervals. Soundappan et al. (2004) present an excellent comparison of evidence theory and Bayesian theory for modeling uncertainty. In their work, they suggest the following assumptions of the evidence theory approach:

1. If some of the evidence is imprecise, uncertainty of an event can be quantified by the maximum and minimum probabilities of that event. In essence these are the absolute bounding probabilities that can be realized from the available evidence.

2. Information about intervals from different experts can be uncertain (i.e. different experts have different opinions) and should be treated as yet another source of randomness and imprecision.

These assumptions relate only to the available data (i.e. bounds) and not to any probability description of the data. This is why some consider the evidence theory approach a lower information augmentation technique to epistemic uncertainty, since no additional information is added by assuming the distribution of the data when only bounding information is given. Be that as it may, the complication in using evidence theory lies in the way uncertainty is represented: by two bounding density functions as

suggested by assumption 1 above. If the uncertainty quantification is aimed at aiding a decision-maker and furthermore the region separated by the two bounding density functions is large, then this is of little help in formulating a decision (Soundappan et al., 2004). Using probability theory to model epistemic uncertainty has its drawbacks but in the context of this research, it provides a feasible approach to incorporating this type of uncertainty into our Bayes network approach. In addition to this, by treating epistemic uncertainty in a probabilistic way, it allows for a single estimate of the probability of the system response quantity of interest (or probability of failure if required) which facilitates the formulation of a risk informed decision making methodology.

Having decided that interval data will be treated in a probabilistic manner, a technique to do this will now be detailed. The approach shown below has been developed at Vanderbilt University by McDonald et al (2009a, b and c) and it is summarized as follows:

1. Obtain from various experts bounding information (i.e. upper and lower bounds on an interval) of a quantity of interest. This can be the interval in which the true parameter of a model might lie, for example the value of damping for a mechanical system.

2. Estimate the bounds on the first four moments of the interval data.

3. Using the range of moments estimated in step 2, establish a family of Johnson distributions using the method of matching moments.

4. Sample the variable from this family of distributions obtained in step 3 and use in a Monte Carlo type analysis for uncertainty propagation.

Because it is impossible to know the true moments of the data given an interval, there are infinitely many Johnson distributions which can represent the interval data. This fact needs to be accommodated in the analysis of interval data. One way to do this is to generate realizations of the moments that satisfy the constraint of falling within the given bounds. From these moments, a distribution from the Johnson family can be obtained. This operation might need to be repeated for several realizations of the moments; this could be done either by a jackknife procedure or a bootstrap procedure (Efron and Tibshirani, 1998) to create realization of the moment. In either case, this implies the need of a double loop in the analysis of uncertainty of the system: an inner loop with the aleatoric uncertainty and an outer loop for the epistemic uncertainty. Details of implementation of this approach will be discussed in Chapter 5.

2.5 Gaussian Process Modeling Overview

Gaussian process (GP) modeling is a technique based on spatial statistics that has been used as a surrogate modeling technique to replace complex and computationally expensive finite element model (FEM) runs, particularly when multiple realizations of the FEM are needed such as in uncertainty quantification and propagation. GP modeling uses a set of observed inputs and outputs (commonly referred to as training data) to construct an approximation to the underlying relationship. It is also desired that the resulting approximation function (i.e. GP model) interpolates within the range of the input data, thus the selection of appropriate training data becomes as important as the creation of the GP model. Since GP modeling is used only as a tool to enable the computations of response quantities, this topic will not be dealt with, here, in great depth. It is briefly

explained for completeness and the interested reader is referred to McFarland, 2008 for an excellent description of the GP algorithm used in this research. (Also refer to Rasmussen, 1996; Martin and Simpson, 2005; Mardia and Marshall, 1984 and Santner et al., 2003 for additional information).

One advantage of a GP model is that it is a non-parametric modeling technique that avoids the need for a functional relationship among data to be known in advance. This provides much flexibility to establishing a relationship between a set of inputs and outputs but it also creates a limitation as to the applicability of such a model. In general, as long as the GP model is used in a regime (i.e. a set of inputs) that does not extend too far outside the boundaries within which the GP model was created, the interpolations will generally be good. A GP model will not normally produce accurate extrapolations. The GP model has another significant feature of interest in this research in that it provides a direct representation of the uncertainty associated with its interpolative approximation. As noted in McFarland, 2008, GP modeling is quite powerful but there is a steep learning curve needed to obtain a working understanding of the methodology, and the implementation can lead to erroneous conclusions if the parameters of the model are not selected carefully. A brief discussion of the theory behind GP models is described next.

Consider that one wants to build an approximation to a function of a vector-valued input X, based only on m observations of the inputs and outputs: $Y(x(1)), \ldots, Y(x(m))$. As noted above, an appropriate selection of these input/output pairs is critical to the performance of the GP model. The basic idea of the GP interpolation model is that the outputs, Y, are modeled as a Gaussian process that is indexed by the inputs, x. A Gaussian process is simply a set of random variables such that any finite subset has a multivariate

Gaussian distribution. A Gaussian process is defined by its mean function and covariance function, which in this case are functions of X. Once the Gaussian process is observed at m locations $x(1), \ldots, x(m)$, the conditional distribution of the process can be computed at any new location, x^*, which provides both an expected value and variance (uncertainty) of the underlying function. The key here is that the function describing the covariance among the outputs, Y, is a function of the inputs, X. The covariance function is constructed such that the covariance between two outputs is large when the corresponding inputs are close together, and the covariance between two outputs is small when the corresponding inputs are far apart. As shown below, the conditional expected value of $Y(x^*)$ is a linear combination of the observed outputs, $Y(x_1), \ldots, Y(x_m)$, in which the weights depend on how close x^* is to each of $x_1, \ldots, x_m$. In addition, the conditional variance (uncertainty) of $Y(x^*)$ is small if x^* is close to the training points and large if it is not. Further, the GP model may incorporate a systematic, parametric trend function whose purpose is to capture large-scale variations. This trend function can be, for example, a linear or quadratic regression of the training points. It turns out that this trend function is actually the (unconditional) mean function of the Gaussian process. The effect of the mean function on predictions that interpolate the training data tends to be small, but when the model is used for extrapolation, the predictions will follow the mean function very closely as soon as the correlations with the training data become negligible.

To develop the theory, let $Y(x)$ denote a Gaussian process with mean and covariance given by

$$E[Y(x)] = f^T(x)\beta$$

and   (2.1)

$$Cov[Y(x), Y(x^*)] = \lambda c(x, x^* | \xi)$$

where $f^T(x)$ defines basis functions for the trend (either, a constant or linear trend); β gives the coefficients of the regression trend; $c(x, x^*|\xi)$ is the correlation between x and x^*; and ξ is the vector of parameters governing the correlation function. Consider that the process has been observed at m locations (the training or design points) $x_1, \ldots, x_m$ of a d-dimensional input variable, yielding the resulting observed random vector, Y. By definition, the joint distribution of Y satisfies:

$$\mathbf{Y} \sim N_m(f^T(x)\beta, \lambda\mathbf{R}) \tag{2.2}$$

where $\mathbf{R}$ is a matrix of correlations among the training points. In this research the parameters of the trend function and the covariance function are estimated using maximum likelihood estimation, thus, the conditional expected value and variance (uncertainty) of the process at an untested location x^* are calculated as:

$$E[Y(x^*)|\mathbf{Y}] = f^T(x^*)\beta + r^T(x^*)\mathbf{R}^{-1}(Y - F\beta) \tag{2.3}$$

and

$$Var[Y(x^*)|\mathbf{Y}] = \lambda(1 - r^T\mathbf{R}^{-1}r) \tag{2.4}$$

where $\mathbf{f}$ is the vector of trend basis functions at each of the training points, and $\mathbf{r}$ is the vector of correlations between x^* and each of the training points.

2.6 Summary

This chapter presented a summary of RIDA and its supporting methodology, QMU. It is noted that QMU will be accomplished by a combination of experimental and computational means. Identifying and quantifying sources of uncertainty are critical steps in this process and the two main categories, aleatoric and epistemic were described. Ideally, experimental data will be used to quantify uncertainty. Due to the complexity of the systems being analyzed, experimental data at the system level is seldom available but data from individual components or group of components is more readily available. This study proposes the use of Bayesian methods to incorporate all available data at different levels of complexity and propagate it through to the system level. In order to quantify the uncertainty at the system level, many function evaluations of the system model need to be performed. Usually, the system level model is in the form of a computationally expensive finite element model. It is proposed in this study that a Gaussian process model, which is a type of surrogate model be used to enable fast and efficient computations of the system level prediction. With this background in place, attention is now turned to the demonstration problem used in this research. This includes a description of the available data and computational models.

CHAPTER III

APPLICATION PROBLEM

3.1 Introduction

Physical systems are complex and the interaction of their parts can result in coupled behaviors. As we seek to model these physical systems we are faced with the need to include different individual mathematical models that represent various types of physics. In this example we look at a model that incorporates both a non-linear mechanical joint model and epoxy-based foam. These models have been developed using constitutive type experiments and formulations and have been independently validated using test data different from that used to calibrate them. The next step was to combine these models and exercise them in an environment different from the environments in which they had been validated. As explained in an earlier chapter, this is a hierarchical approach to building complex system models. The example uses an aerospace component developed at Sandia National Laboratories and the details of the experiments and modeling of the components are described in the following sections.

At Sandia, there has been an emphasis on developing models of components from first principles, calibrating them from simple exploratory experiments, validating them relative to a different set of experiments and then using them within a more complex model. For example, one could investigate the behavior of a mechanical joint, develop a model that explains some phenomenon, validate its performance (based on a different use environment) and use it as part of larger system. Likewise, one could develop another

31

model for the behavior of encapsulating foam and repeat the same sequence of steps as for the mechanical joint. Similar steps could be taken for various other components which then will aggregate to form a full system (i.e. a system which is composed of joints, foam and other components). In addition, there can be multiple tests of these components and thus a probabilistic analysis of the data could be made. Furthermore, it is possible that the interactions of the various components were never tested, thus no information on the coupling of components is available or interactions of components were tested except at excitation levels that are not comparable to those of the full system. What was described above is a hierarchical approach to building a system level model and has been repeatedly used at Sandia and other Department of Energy research laboratories. It basically is a construction of a complex system model by using a building block approach that incorporates simpler component based models and couples them together. This hierarchical model building approach was described in Oberkampf et al. (2000), Sindir et al. (1996) and it has been implemented by Urbina et al. (2006) for the problem being examined in this dissertation.

3.2 Overall Problem Description

The example problem shown in Figure 3.1 was chosen to implement the ideas developed in this research which addresses uncertainty quantification and propagation in a hierarchical model development.
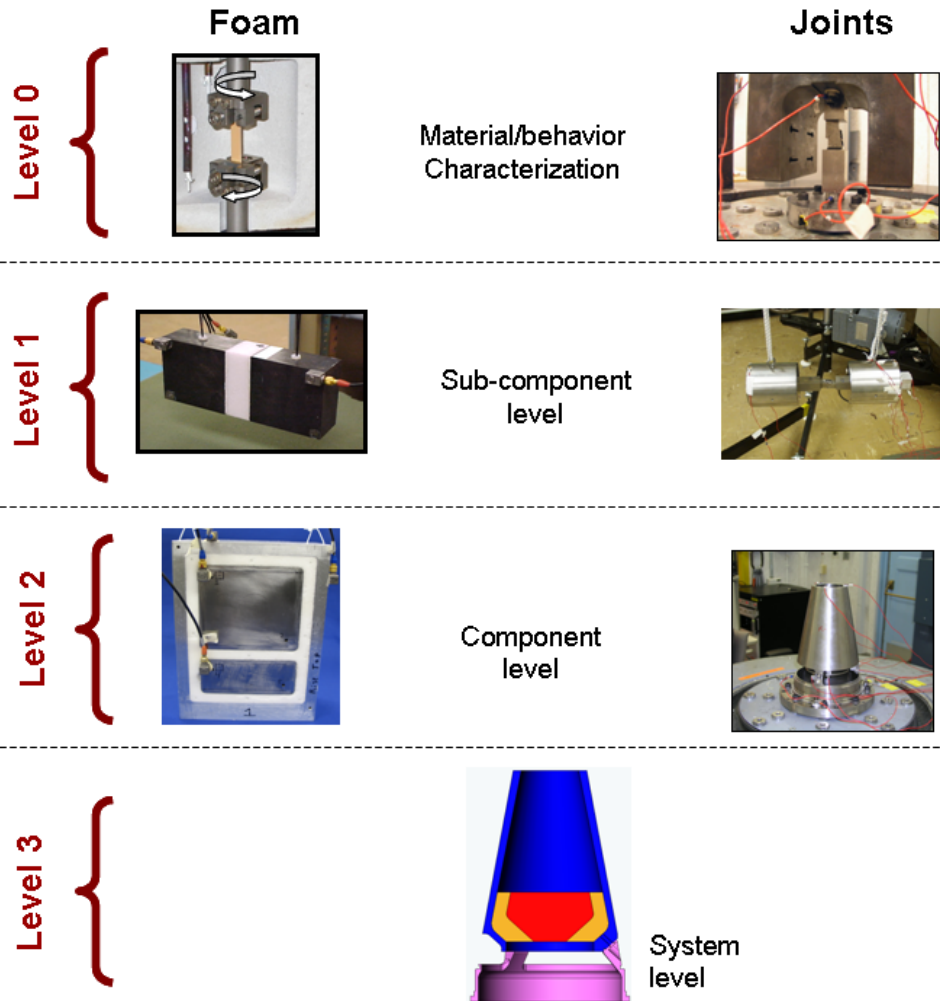
Figure 3.1. Example problem for uncertainty quantification analysis

This problem was originally developed at Sandia for the purpose of implementing the rigorous uncertainty quantification and model validation methodology developed there (Alvin, et al., 2000, Trucano, et al., 2001 and Pilch, et al., 2001). The problem has the following features:

1. It is a 2-component problem where some components are mechanical joints and another component is epoxy-based foam. Both are energy dissipating mechanisms.

2. It is a multi-level problem where the phenomena observed at the lowest levels are assumed to be present at subsequent levels, i.e. damping behavior in the joints at level 1 is assumed the same at levels 2 and 3. A priori, this seems like a reasonable assumption to make but it might turn out to be an incorrect one.

3. The individual component branches converge to a system level hardware where the two couple together. The coupling of the two branches is a potential source of epistemic uncertainty since this interaction had not been tested previously and thus no data is available.

4. Experimental data consists of repeated tests on several nominally identical hardware components, at levels 1 and 2 in Figure 3.1. These are intended to quantify the variability inherent in a physical system due to manufacturing variations as well as test to test variability due to slight changes in the test configuration.

5. Finite element models are built and calibrated to simulate a particular behavior of the physical hardware, at levels 1 and 2. The model parameters have been calibrated from simple, discovery experiments (shown as level 0) aimed at isolating the particular physical phenomenon that the model is meant to represent.

Several observations regarding the example problem and their potential implications are:

1. In general, there is no one-to-one correspondence between the hardware that was tested at the various levels. For example, the single joint tested in level 1 is not part of the 3-leg system at level 2. Similarly for the foam, the piece of foam in level 1 is not the same one (nor does it come from the same batch) as the foam in level 2. This issue could make relating data from one level to another difficult.

2. The type of data collected for the joints and foams are different. In general, for the joints, time domain data is available and for the foam, frequency domain data is available.

3. Model runs can be made if necessary, at any level. Surrogate models, particularly of the system level, might be necessary to expedite the calculations.

4. No system-level experimental data is available. This makes all the predictions at the system level extrapolations.

The data available for this example are described below.

3.3 Available Experimental Data and Models

In this section, the experimental data and models available for this example at levels 0, 1 and 2 and the model at full system level (no data available) are described. It is assumed that all experimental data collected has been quality checked and does not contain any systematic errors. In addition, it is assumed in this research that measurement error is negligible relative to the sources of uncertainty examined in this research. Also in this section, the computational models are described as well as the Gaussian process (GP)

models used in lieu of the finite element models. Use of GP models are needed to make the uncertainty quantification and propagation a computationally tractable problem.

3.3.1   Foam - Level 0: Material Parameter Characterization

The first step in the process is to characterize the material model that describes the behavior of epoxy-based foam with a nominal density of 20 pcf (pounds per cubic foot) and also to quantify the variability in the material. Samples used for this characterization are shown in Figure 3.2.
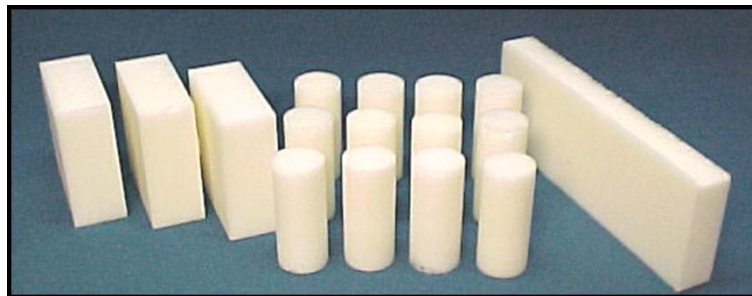
Figure 3.2. Samples for level 0 experiments

The material density of the foam samples was estimated through physical measurements and the elastic and shear modulus were estimated using tension/compression and torsional experiments. The experiments and resulting data are shown in Figure 3.3 and Figure 3.4.

36

Figure 3.3. Torsional and tension experiments used for calibration
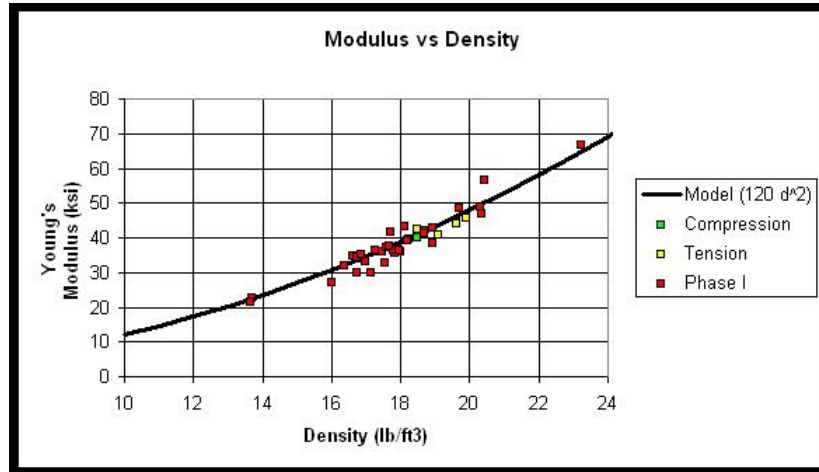
Figure 3.4. Calibration data for foam modulus of elasticity

Based on the data shown in Figure 3.4, a model describing the relationship between density and modulus of elasticity was chosen from Gibson and Ashby, 1999. The model is

$$E = c_1 * (\rho)^{pow} \tag{3.1}$$

37

The model parameters c_1 and *pow* were calibrated through regression using the data in Figure 3.4, and estimated as $c_1 = 120$ and $pow = 2.0$. In addition to calibrating a model of foam modulus of elasticity, the data in Figure 3.4 is used to obtain statistics (mean and standard deviation) of the foam density.

3.3.2 Joints - Level 0: Material Parameter Characterization

For the joints, experiments consisting of sine sweeps at 5 load levels (100, 200, 300, 400 and 500 lbs) were performed and data in the form of force versus energy dissipation were used to calibrate a constitutive model that represents the behavior of the physical joint. Variability data from sample to sample and test to test were also obtained from a total of 9 combinations of three top and three bottom pieces of a bolted connection. These samples are shown in Figure 3.5.

Figure 3.5. Single leg samples for calibration experiments

One of the combinations of the joint is shown in Figure 3.6 while mounted on an electro-magnetic shaker. Each of the nine hardware combinations was assembled and disassembled 5 times in order to quantify test to test variability as well as sample to sample variability.
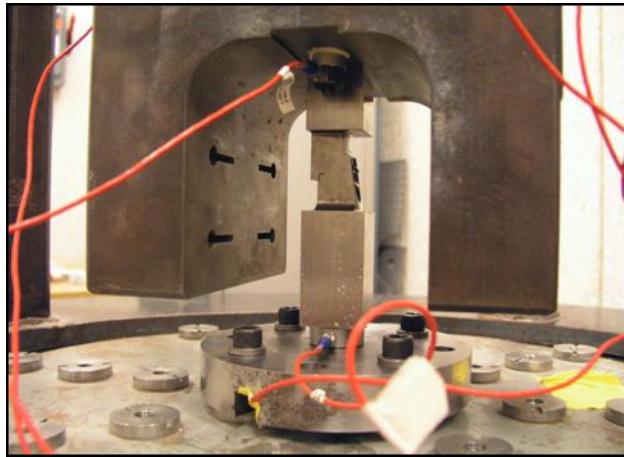
Figure 3.6. Single leg test fixture

Starting from acceleration time histories recorded from each experiment and using a log decrement approach to estimate the damping of the system, a collection of curves relating force versus energy dissipation were calculated. These are shown in Figure 3.7.
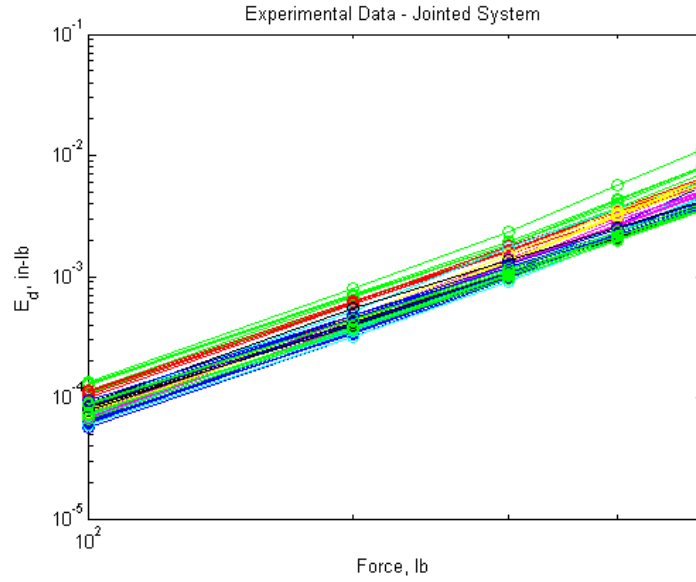
Figure 3.7. 45 realizations of force vs. energy dissipation curves

In addition to this data, information regarding the hysteretic behavior of the joint is also obtained from the experiments. A schematic of this type of data along with some salient features is shown in Figure 3.8.
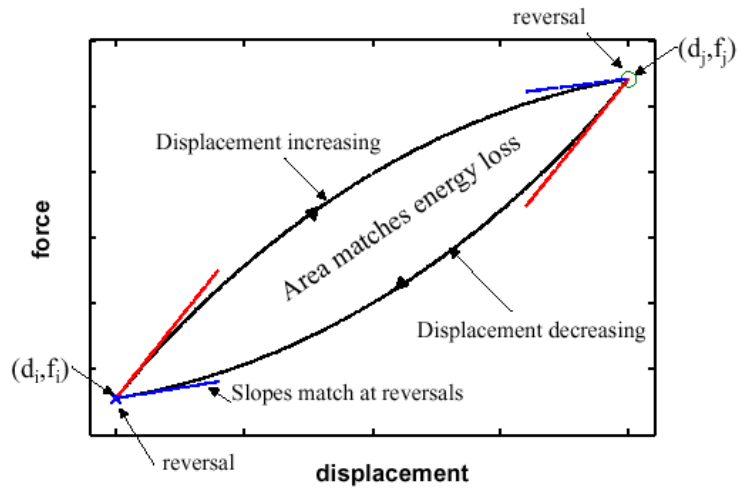
Figure 3.8. Hysteresis curve for joint behavior

40

Using the data shown in Figure 3.7 and Figure 3.8, a constitutive model relating the energy dissipated per cycle of response to force across the joint can be calibrated. In this case, the following fmodel proposed by Smallwood (Smallwood et al., 2000) was used:

$$f_j = k_{lin}(d_j - d_i) - k_{non}(d_j - d_i)^{npow} + f_i \qquad (3.2)$$

where

$f_{i\&j}$ = force across the joint at the reversal point in the hysteresis loop (measured during experiment)

$d_{i\&j}$ = relative displacement across the faces of the joint at the reversal point in the hysteresis loop (measured during experiment)

k_{lin} = linear stiffness component. This is a calibrated parameter.

k_{non} = non-linear stiffness component. This is calibrated from the data.

$npow$ = degree of nonlinearity. It is the slope of each experimental curve of E_d vs. Force curve shown in Figure 3.7 in log-log space

Since there are 45 repeated experiments (9 hardware combinations times 5 assembly/disassembly combinations), a probability distribution for each of the model parameters can be created.

3.3.3   Foam - Level 1: Component Level

Tests at this level targeted the stiffness properties of the epoxy-based foam by placing it in a 2 degree of freedom configuration. Since the foam is bonded to the steel

41

masses, these tests also examined the foam's damping properties. The structure is an element of foam (shown as a white block in the center of the test hardware) bonded to steel end masses for which six samples were available. Modal tests were performed on the samples with a representative configuration shown in Figure 3.9 and a schematic of the first 4 modes of vibration are shown in Figure 3.10.
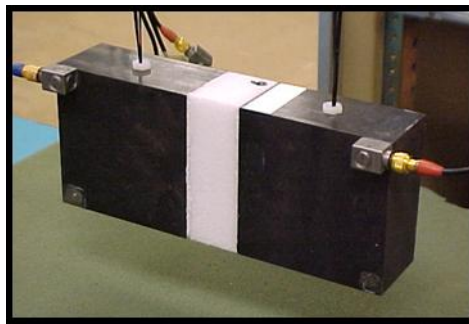
Figure 3.9. Modal testing setup for foam level 1

1st Bending X

Torsion

1st Bending Y
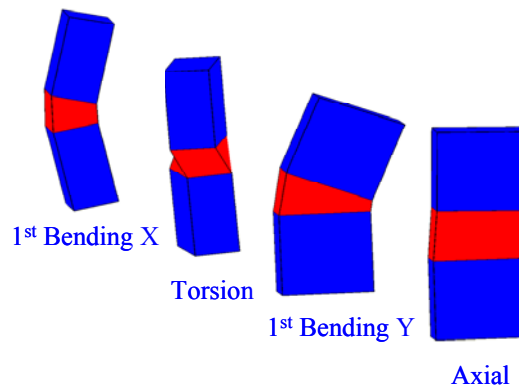
Axial

Figure 3.10. First 4 bending modes for foam level 1 sample

Based on the anticipated behavior of the system, the natural frequency in the axial direction was selected as the response measure that was most relevant to the system level behavior.

A finite element model of the test hardware shown in Figure 3.9 was created and is shown in Figure 3.11. The model consists of 1470 Hex8 elements and 1920 nodes and it was analyzed in Salinas, a linear structural analysis code (Reese, 2004). Convergence studies were performed using four different mesh sizes to show model verificaton. Eigenvalues were used as the convergence metric and the Richardson extrapolated convergence errors (Richardson, 1910) were found to be less than 1.3% for the first six natural frequencies of the flexible modes of vibration in the mesh chosen. Modal analyses were performed and natural frequencies were obtained. The natural frequency in the axial direction will be used in this study.
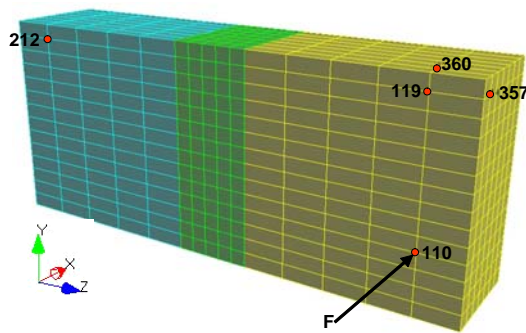
Figure 3.11. Finite element model of foam hardware for level 1

3.3.4 Joints - Level 1: Component Level

Similarly to the foam level 1 experiments, tests at this level targeted the stiffness and energy dissipating properties of the joints while in a 2 degree of freedom configuration. This configuration places the single leg joint hardware used for the behavior characterization done in level 0 in a different loading configuration. Two 30 lb masses are bolted at the ends of the single leg creating a "dumbbell" shaped hardware. This is shown in Figure 3.12.
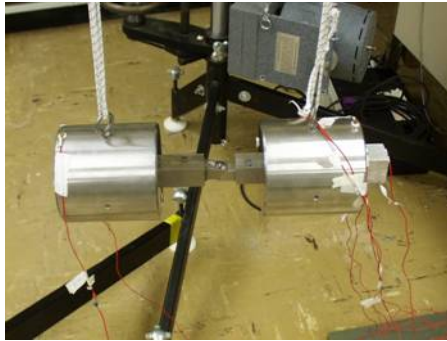
Figure 3.12. One configuration of the single leg joints for level 1 experiments

This configuration is then supported by bungee cords to simulate a free-free environment and it is subjected to an impulse excitation provided by an instrumented hammer. The acceleration response of the dumbbell on the end opposite to the excitation end is recorded and multiplied by the mass of one of the dumbbells to obtain the force across the joint. This is shown in Figure 3.13.
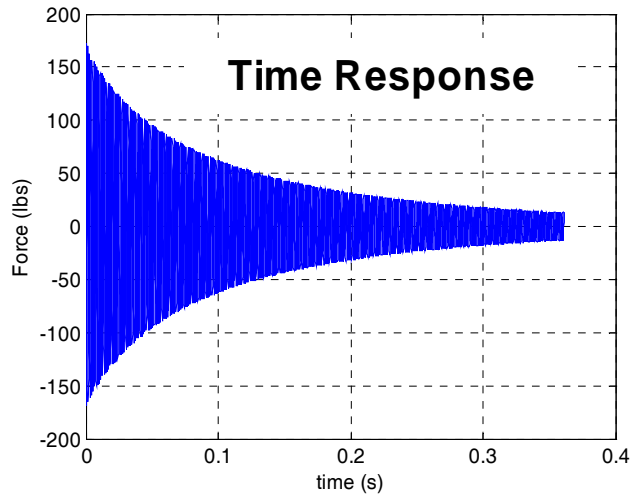
Figure 3.13. Time history of force across the joint (from F=ma)

From this, an estimate the energy dissipation of the system at a particular force level can be obtained. This is done with the formulation shown below:
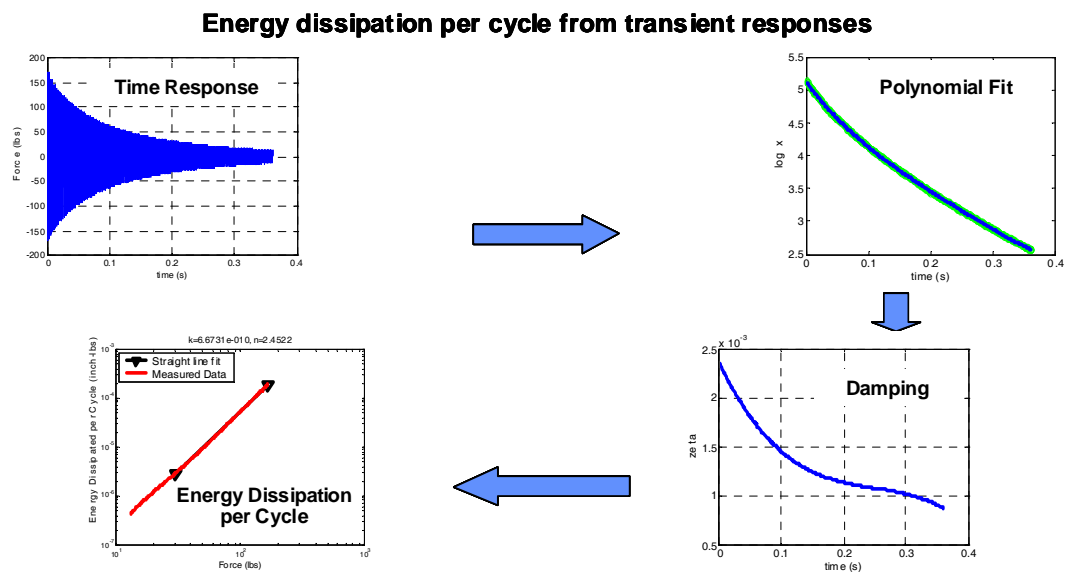
Figure 3.14. Schematic of the process to calculate energy dissipation per cycle from transient responses

To implement the formulations shown in Figure 3.14, start with a time response that is in free decay and write the expression:

$$x(t) = e^{-\xi \varpi_n t} \cos(\varpi_d t) \tag{3.3}$$

where ξ, ϖ_d and t are the damping ratio, damped frequency and time, respectively and x is the measured signal. From this define the envelope of the peaks of the signal defined in Eq. 3.3 as:

$$x(t) = e^{-\xi \varpi_n t} \tag{3.4}$$

and taking the logarithm of Eq. 3.4 and subsequently the time derivative to get:

$$\log(x(t)) = -\xi \varpi_n t$$

$$\frac{d(x(t))}{dt} = -\xi \varpi_n \tag{3.5 and 3.6}$$

Finally to estimate the energy dissipated per cycle of response, E_d, at a given force level, F, use the following expression:

$$E_d = \frac{\xi F^2}{m^2 f_n^2} \tag{3.7}$$

where m is the mass and f_n is the natural frequency of the system.

A total of 45 experiments were conducted which consisted of 5 repetitions of the experiment for each of the nine leg configurations and the resulting energy dissipation curves are shown in Figure 3.15.
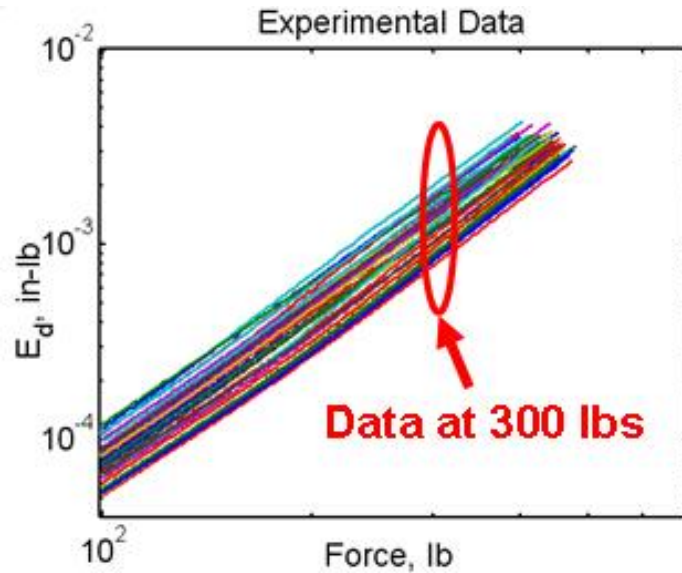
Figure 3.15. Energy dissipation vs. force curves for level 1 tests

From this data, it was determined that the energy dissipation at 300 lbs was the most relevant at the system level and thus chosen as the response measure of interest.

A finite element model for the hardware shown in Figure 3.12 was simplified to a few-degrees-of-freedom model that has all the mass properties of the actual hardware and the nonlinear joint element. This permits fast analysis of the model when time domain computations must be performed. The schematic of the model is shown in Figure 3.16. It consists of 4 nodes and 3 elements, one of which is the Smallwood element. This model was analyzed using Salinas. Due to a limitation of Salinas, accelerations cannot be used, directly, to excite the model so the experimental acceleration excitation was converted to a force.
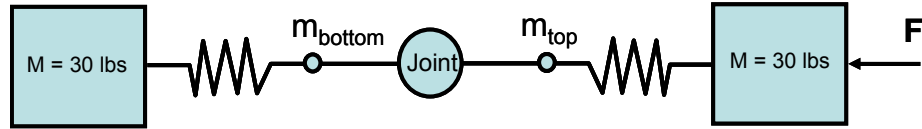
Figure 3.16. Lumped mass model of joints at level 1

3.3.5   Foam - Level 2: Sub-system Level

This configuration starts to simulate the conditions that will be present at the system level. Namely, that the foam encapsulates a rigid component and it transmits and dampens externally applied loads relative to the encapsulated mass. Six different test specimens were fabricated in order to investigate part to part variability. These are shown in Figure 3.17.
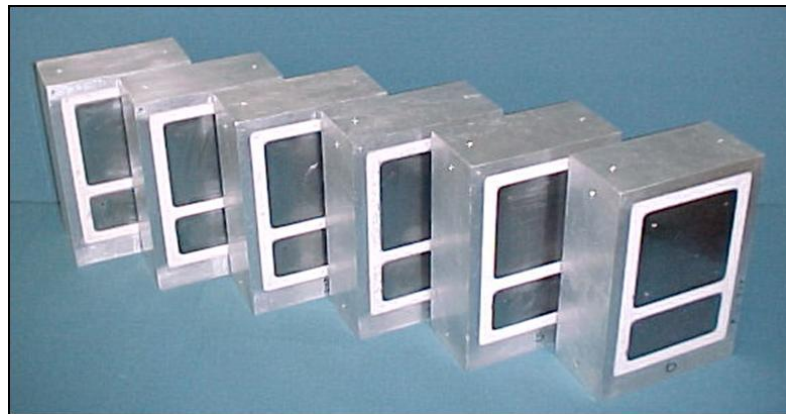
Figure 3.17. Foam level 2 samples for modal testing

This hardware consists of a set of steel masses encapsulated in foam and contained within an aluminum outer shell. The hardware was instrumented with four triaxial

accelerometers in order to minimize mass loading effects and additional damping introduced by the cables. The instrumented hardware is shown in Figure 3.18.
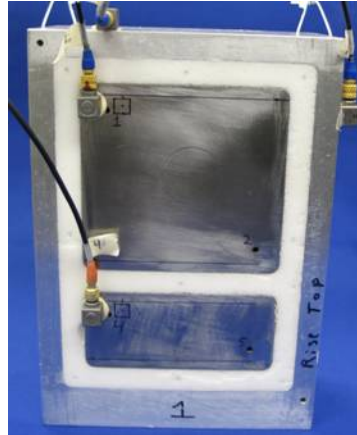
Figure 3.18. Foam sub-system test article

The test fixture was suspended from bungee cords to simulate free-free conditions and excited using a small, instrumented hammer. All testing was done at room temperature and acceleration time histories at each accelerometer location were recorded and mode frequencies and mode shapes were extracted. Of special interest is the natural frequency corresponding to the axial mode of vibration which is shown schematically in Figure 3.19.
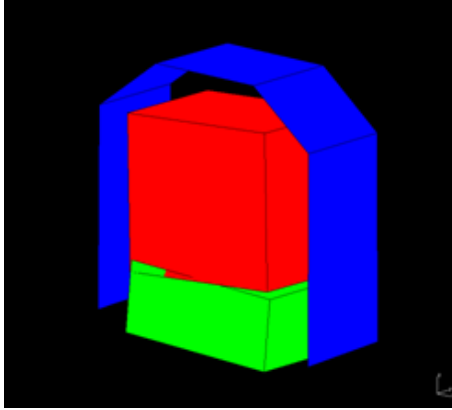
Figure 3.19. Axial mode of foam level 2 hardware

A finite element model (FEM) of the hardware shown in Figure 3.17 was constructed and used to simulate the behavior of the test specimen. A schematic of the FEM is shown in Figure 3.20.
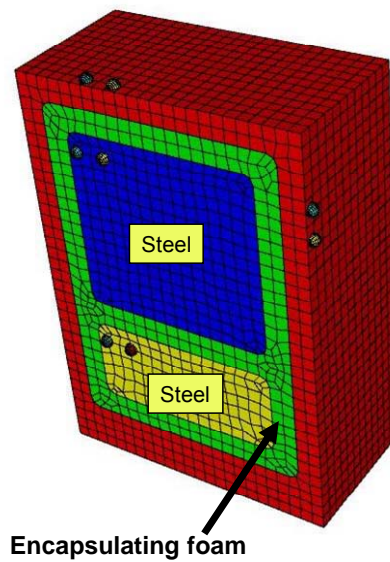
Figure 3.20. Foam level 2 finite element model

The model is comprised of approximately 12000 nodes, 10000 eight node hex elements and a modal analysis was performed. A perfect bond was assumed between the foam, the aluminum outer shell and the steel masses, and furthermore, the foam was assumed to be perfectly homogenous (i.e. no voids or substantial changes in density). The foam was modeled using a linear-elastic type formulation with parameters derived from simpler constitutive tests. Mesh convergence studies were performed to assess the suitability of the proposed mesh discretization and it was determined through Richardson extrapolation analysis that the error due to mesh size was very small relative to the uncertainty present in the experimental data. Analysis of this model was performed using in Salinas.

The functional relationship to be modeled with a surrogate model is the one between modulus of elasticity and the axial mode for the level 1 and level 2 hardware. Realizations of the functional relationship between modulus of elasticity and axial mode frequency are shown in Figure 3.21 for the Level 1 and Level 2 structures.
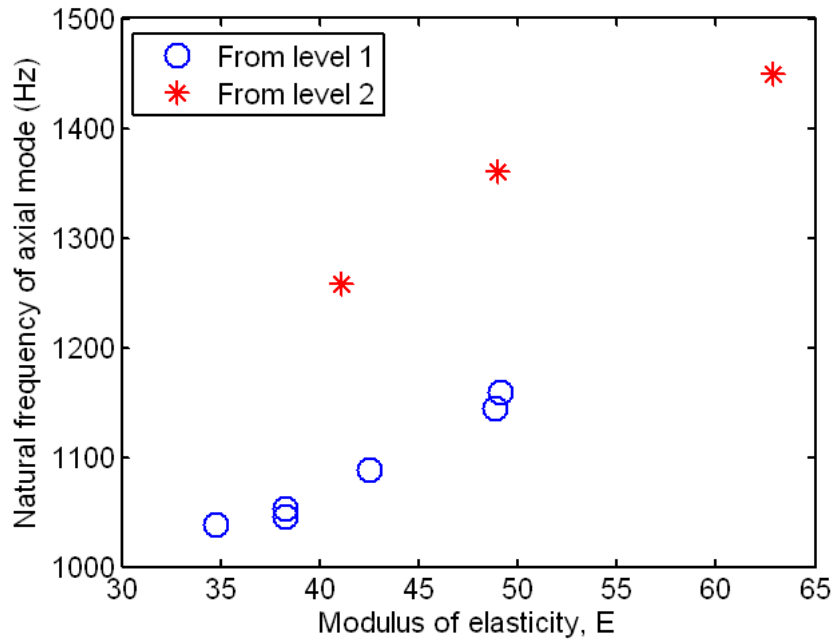
Figure 3.21. Modulus of elasticity versus natural frequency in the axial direction

As it can be observed, the relationship, given the small number of test data, is a simple one and thus will be modeled with a polynomial fit to the data. The forms of these models are:

$$f_n = aE^2 + bE + c \tag{3.8}$$

where f_n are the natural frequency of the axial mode of vibration for either the level 1 or level 2 hardware, E is the modulus of elasticity and a, b and c are coefficients of the polynomial model. The estimated values of parameters for each level are shown in the table below:

Table 3.1. Coefficient for polynomial fit of natural frequencies as a function of modulus of elasticity for foam level 1 and 2

Level	a	b	c
1	0.3239	-18.891	1301
2	-0.3023	40.215	115

3.3.6   Joints - Level 2: Sub-system Level

The hardware and the tests at this level closely approximates the final joint configuration at the system level. The stiffness and damping properties tested at this level could be a good starting point to estimate those at the system level. The experimental system is a truncated conic shell supported on legs at three approximately equidistant locations (around the curcumference). The support structure beneath the legs is a cylindrical shell – relatively thin on its top, and transitioning into a thicker section. The conic shell is attached to the support structure via three screws, each of which passes through a hole in a thin, flat plate at the top of a leg. Three nominally identical replicates of the conic shell were fabricated, along with three nominally identical support structures. They are shown in Figure 3.22.

Figure 3.22. Test articles for joints level 2

The nine combinations of shells and support structures were tested in environments generated using an electrodynamic shaker in the laboratory. The holes in the base of the support structure were used to attach it to an adaptor plate that was connected to the shaker armature. The test setup is shown in Figure 3.23.
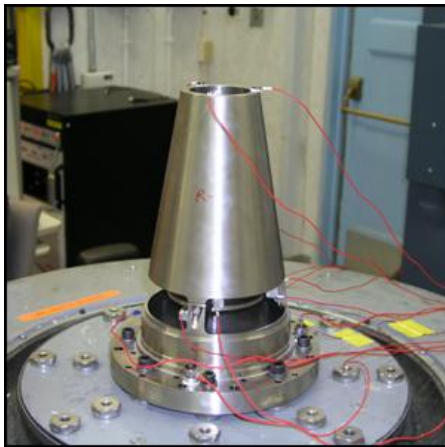
Figure 3.23. Shaker testing for joints level 2

Each of the nine test structures was excited by a wavelet type input. Each shell-base combination was assembled, disassembled and reassembled three times, and tested each time. The average acceleration structural responses at the tops of the encapsulated masses were recorded and yielded twenty-seven time histories – nine structures times three tests each. These responses were then used to estimate the energy dissipation per cycle of response at various force levels using the formulation shown in Figure 3.14 and Equations 3.3 through 3.7. For this research, the energy dissipated corresponding roughly to 300 lbs of force across each leg was selected as the response measure of interest. The collection of energy dissipation curves vs. force is shown in Figure 3.24.
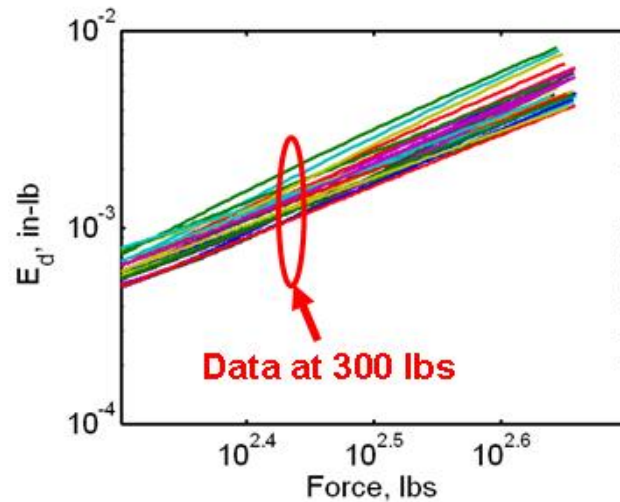
Figure 3.24. Energy dissipation vs. force for joints level 2

The model for the physical system shown in Figure 3.23 is simply the lumped-mass representation shown in Figure 3.25.
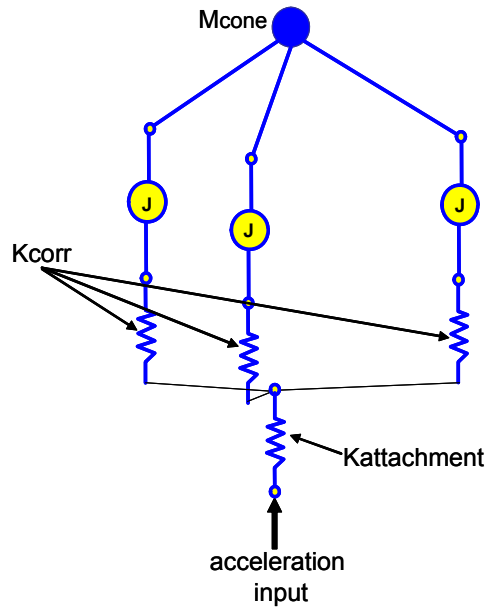
Figure 3.25. Lumped model of joints level 2

The element thought to be critical in the model is a nonlinear spring (denoted J in the figure), and it is modeled using the framework of the so-called Smallwood element. The parameters of the nonlinear, Smallwood spring element were identified based on experiments in which individual leg-simulators were excited sinusoidally. Multiple systems were tested, and they are stochastic, therefore, the parameters of the Smallwood model are described in a probabilistic framework. Because the geometry and boundary conditions of the system used to identify Smallwood model parameters differ from the geometry and boundary conditions of the three-legged system, a correction stiffness, K_{corr} and an attachment stiffness, $K_{attachment}$ must be inserted into the lumped mass model to render its predictions accurate. The attachment stiffness was calibrated by matching the axial frequency of a monolithic structure and assuming that the stiffness of the cone is

rigid when compared to the rest of the structure. The correction stiffness was calculated and inserted into the lumped-mass model, and predictions of the system acceleration response were made. Analysis of this model was done in Salinas. Each model was excited with the input waveform resembling a wavelet which has similar dynamic characteristics as those used in the experiments. Acceleration time histories for each model prediction were obtained.

For both level 1 and 2 models, Gaussian process models (GPM) were developed that relate the linear and non-linear stiffness components and the degree of non-linearity to the energy dissipated by the joint at an input level of 300 lbs. This model was trained with approximately 75% of the available full model simulations and the remainder were used for testing the GPM. The test data was used to calculate the root mean square error and it was determined to be 6.78e-5 while the average error between the predicted and the test data is less than 3% for both levels 1 and 2 surrogate models. It was then deemed that the GP models were good representations of the full scale model which relate the three joint parameters (k_{lin}, k_{non} and npow) to the energy dissipated at 300 lbs.

3.3.7   Level 3: System Level

There is no experimental data at this level. A finite element model of the physical system was constructed using the Sandia-developed CUBIT meshing tool (CUBIT, 2008) and analyzed using Salinas. Figure 3.26 shows the exterior and interior of the system model and the input and output locations.
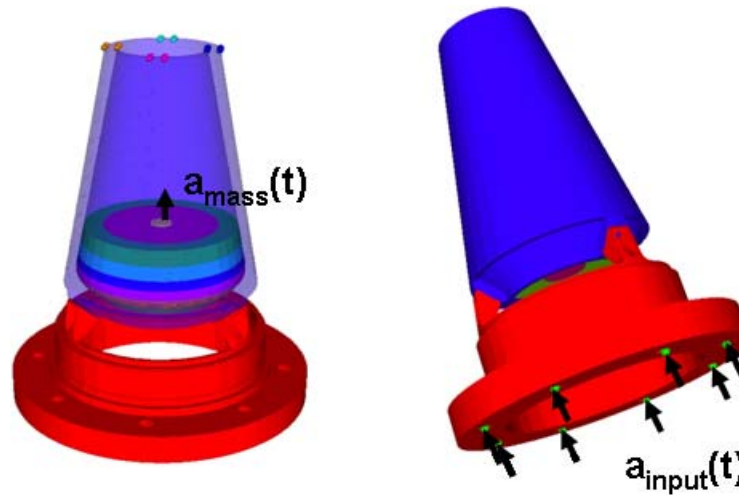
Figure 3.26. Finite element model with input and response locations

The finite element model uses the Smallwood model in Equation 3.2 to represent the nonlinear energy dissipation behavior of the bolted connection between the conic part and the lower assembly, and the linear-elastic model for the encapsulating foam component. All solid pieces [the conic part, the bottom piece and the internal encapsulated mass (shown in the cross section in magenta color)] are made of stainless steel. The encapsulating foam is shown as various colored layers and uses the same type of foam which was used in level 0, 1 and 2. Full adhesion is assumed between the epoxy foam and the inside of the conic section and between the epoxy foam and the encapsulated mass. The model consists of 8052, 20-node, hexagonal-type elements which yields approximately 42,000 nodes in the model and was verified by doing a Richardson extrapolation on the natural frequencies. Multiple non-linear, transient analyses were then performed to predict structural response. When subjected to a blast type input the

acceleration time history response at the top of the encapsulated mass is shown in Figure 3.27.
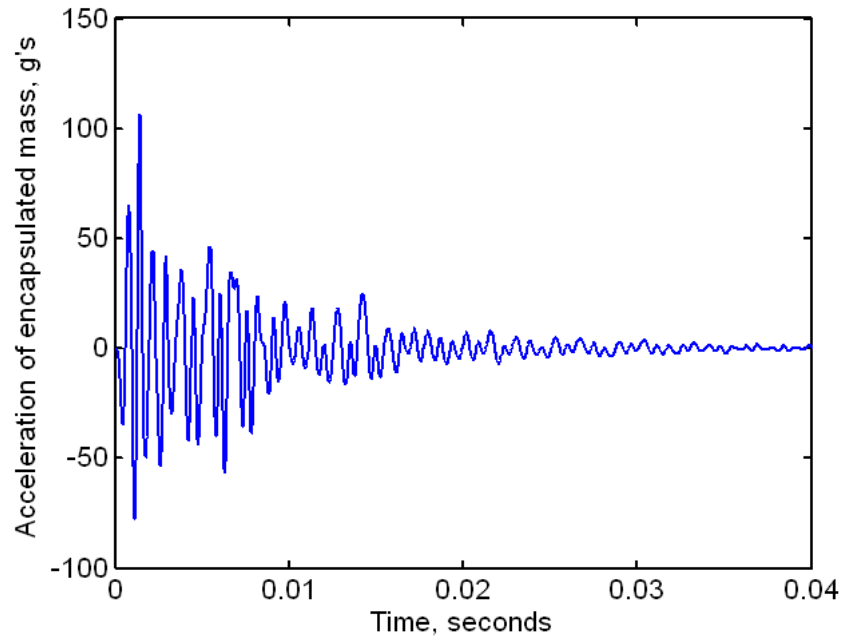
Figure 3.27. Model predicted acceleration response at top of encapsulated mass

These multiple transient analyses form the dataset available to build and test a GP model to capture the relationship between the input (1) linear stiffness, (2) nonlinear stiffness, (3) degree of nonlinearity and (4) modulus of elasticity to the output, absolute peak acceleration of the encapsulated mass. Since the main focus of this research is the modulus of elasticity (as the source of epistemic uncertainty), the functional relationship between modulus and absolute peak acceleration is shown below.
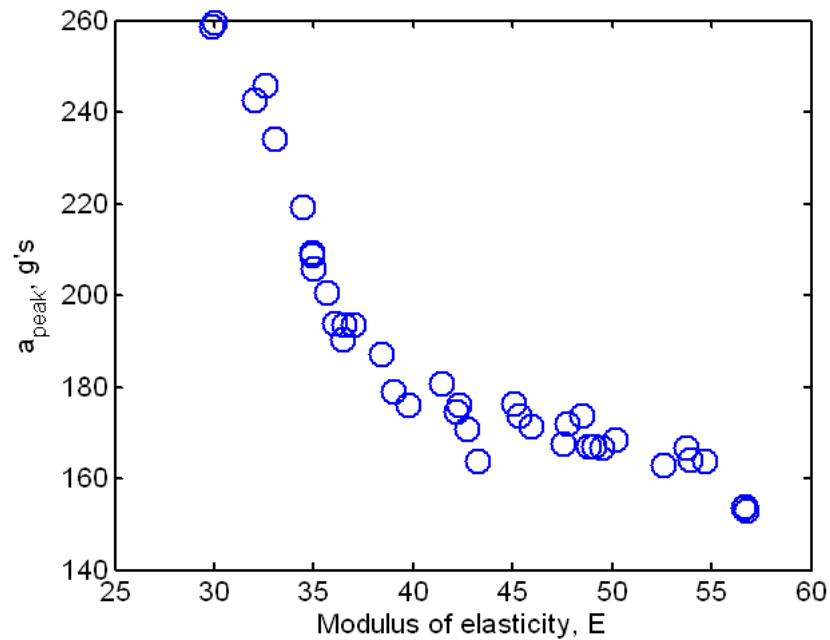
Figure 3.28. Modulus of elasticity versus absolute peak acceleration at encapsulated mass

An important observation involves the inverse relationship between modulus of elasticity and peak acceleration. This relation was at first counter-intuitive; it might be anticipated that an increasing modulus of elasticity (or system stiffness) should cause a corresponding increase in the peak acceleration. In this case, it does not, because an increase in the modulus of elasticity increases the center of a frequency band over which substantial response amplification occurs for the structure. Because the excitation has fixed signal content, such a change in the center frequency can move the region of high amplification to a range where lower input signal content exists. That is the case for this input and this structure. This is an important issue that points to the need to fully understand the system being evaluated. It also serves as a reminder that, in a very complex system, multiple factors can contribute to counter-intuitive results.

A Gaussian process (GP) model was developed that relates the linear and non-linear stiffness components, the degree of non-linearity in the joints and the modulus of elasticity of the foam to the absolute peak acceleration of the encapsulated mass. This model was trained with approximately 60% of the available full model simulations and the remainder were used for testing the GP model. The test data was used to calculate average error between the predicted outputs and the measured test outputs and it is less than 2%. At this point, the results from the GP model are consider consistent with the full finite element model and will be used for the uncertainty quantification and propagation component of this research.

3.4 Summary

The experimental data and the simulation models available for this study have been presented in this chapter. It is easy to see that there is a wealth of data available to use in this example and thus is imperative to condense it into a useable form and better yet, into a form that reflects the main characteristics of the system that is ultimately being assessed. The data at the component and sub-system levels are relevant at the system level and can be used to update the parameters of the system level model in order to make the best prediction at the system level. Using the available data, we will concentrate on quantifying and propagating both aleatoric and epistemic uncertainty to quantify the uncertainty in system-level response prediction in the next two chapters.

CHAPTER IV

PROPAGATION OF ALEATORIC UNCERTAINTY IN A HIERARCHICAL MODELING

4.1 Overview

In this chapter it is proposed to use a Bayes network as a tool to integrate observed data and prior knowledge within a hierarchically built system level model. Bayesian updating propagates uncertainty through this network up to the system level response of interest. In this chapter, the methodology shown in Rebba and Mahadevan (2006) will be applied to the demonstration problem described in Chapter 3.

4.2 Bayes Network

To briefly summarize the viability of using Bayes networks for this problem, consider the basic features of a Bayes network. The main reference for this discussion is Jensen (2001). The purpose of a Bayes network is to use statistical and functional relationships among the variables involved in the model to propagate updated information from one variable to another, based on new data. A Bayes network consists of the following:

- A set of variables and a set of directed edges (arcs) between variables

- Each variable has a finite set of mutually exclusive states

- The variables together with the directed edges form a directed acyclic graph (DAG). DAG's do not allow circular causality.

- To each variable B with parents $A_1, \ldots A_n$, there is an associated conditional probability $P(B|A_1, \ldots A_n)$.

In the intended use of a Bayes network, one would like to calculate the posterior probability density function (PDF) associated with some nodes of interest in the network. To do this, it is known that each parent node has a PDF associated with it and each child node has a conditional probability density function, given the value of the parent node. The entire network can be represented using a joint probability density function which is given by the general expression:

$$P(U) = \prod_i P(X_i \mid parent(X_i)) \qquad (4.1)$$

where $P(U)$ is the joint probability of the network and X_i is the i^{th} node in the network. The Bayes network also facilitates the inclusion of new nodes that represent the observed data and thus the updated densities can be obtained for all the nodes. The joint probability density function for the network can be updated using Bayes theorem when data is available. The expression for Bayes theorem is:

$$f_\theta(\theta \mid Y) = \frac{f_\theta(\theta)f_\theta[Y \mid X(\theta)]}{\int f_\theta(\theta)f_\theta[Y \mid X(\theta)]d\theta} \qquad (4.2)$$

Equation 4.2 can be implemented using Markov Chain Monte Carlo techniques (Gilks *et al*, 1996). The marginal PDF of any node in the Bayes network can be obtained by the integration of the joint PDF over all the values of the remaining variables. Thus the Bayes network approach offers methodology to extrapolate inferences from component level information to the system level, as long as the two levels have common, linking node and

the physics does not change. With this background now we look at the specific implementation of the Bayes networks to the example problem.

The Bayes network constructed to address the example problem used in this study is presented in Figure 4.1.
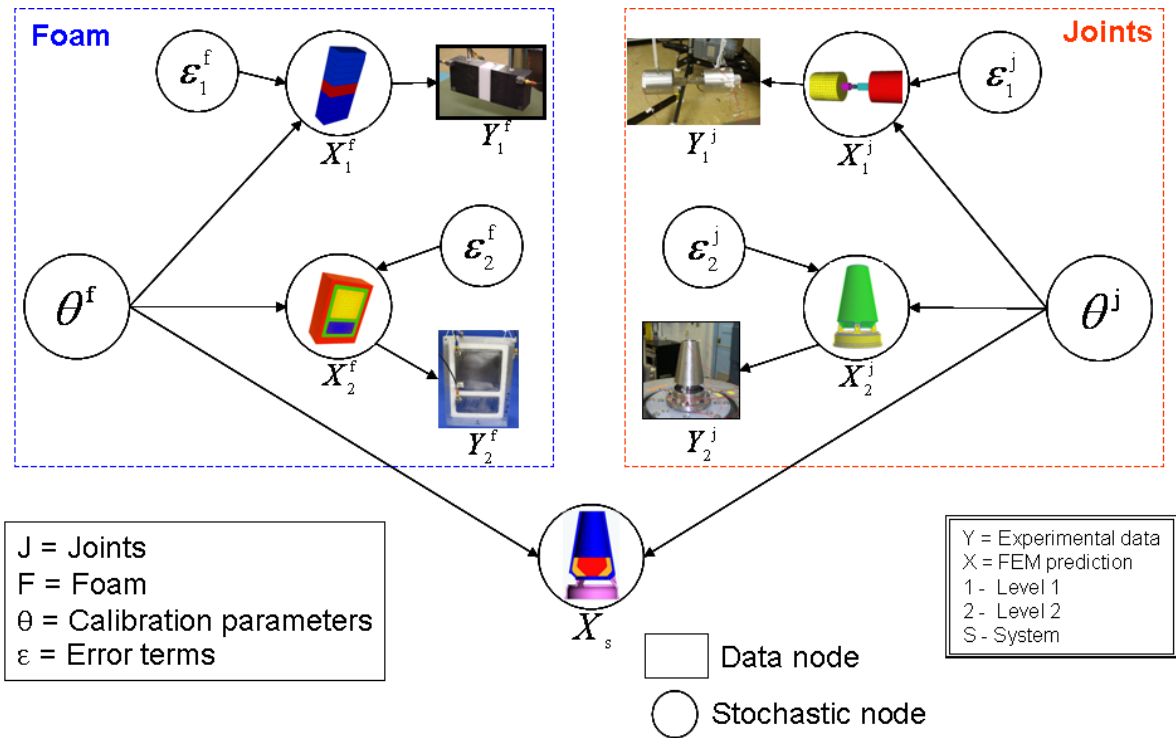
Figure 4.1. Bayes network representation of example problem

From Figure 4.1, one formulates the joint probability density function of the entire network, U as:

$$f(U) = f(\theta^f) * f(X_1^f \mid \theta^f) * f(Y_1^f \mid X_1^f) * f(X_1^f \mid \varepsilon_1^f) *$$
$$f(X_2^f \mid \theta^f) * f(Y_2^f \mid X_2^f) * f(X_2^f \mid \varepsilon_2^f) *$$
$$f(\theta^j) * f(X_1^j \mid \theta^j) * f(Y_1^j \mid X_1^j) * f(X_1^j \mid \varepsilon_1^j) * \qquad (4.3)$$
$$f(X_2^j \mid \theta^j) * f(Y_2^j \mid X_2^j) * f(X_2^j \mid \varepsilon_2^j) *$$
$$f(X_s \mid \theta^f, \theta^j)$$

Note that the error nodes ε shown in Figure 4.1 are associated with the discrepancy in the model predictions relative to the observed data at each level. This error term is assumed to have a normal distribution with zero mean and some standard deviation, σ which will be updated. The main reason behind this is in the way the network is set up. Note that there is a conditional probability that relates the model at each level back to a set of model parameters denoted θ. These model parameters are updated based on the observed data at various levels. It is expected that the models and their parameters can only account for some of the behavior resulting in the observed data (i.e. no model is a perfect representation of the underlying phenomenon). It is generally agreed that the difference can come from various sources, such as model representation of the actual hardware (i.e. dimensions, boundary conditions, contact surface areas, etc), the inability to exactly represent mathematically the dominant physical phenomenon, and model convergence issues (such as solution approximation). To account for these differences, an error term is included in the formulation following, Kennedy and O'Hagan, (2000), McFarland (2008), Landes et al (2006) and Williams et al (2006).

This error term contains among other sources, model form errors and solution approximation errors. These errors are summarized in Rebba and Mahadevan (2006) and methodologies are also suggested for treating them. For the purpose of this research, the solution approximation errors which include mesh discretization and solution stability will be assumed to be addressed by code verification activities which address the

question: Are the numerical solutions correctly implemented? This means, in the context of this research, that issues such as mesh convergence have been addressed and it has been determined that those errors are small relative to other sources of uncertainty such as model form error. Model error will be considered in this research and will be considered a reducible source of uncertainty since it can be reduced if the "correct" model is chosen. This type of error can be quantified when observed data and model predictions are available. In the context of Bayes networks, this is accounted for as a node in the network and is related to the standard deviation of the error in the posterior distribution of the model's prediction as shown below:

$$Y = X + \varepsilon \tag{4.4}$$

where X and Y are the model predictions and the observed data, respectively, at a given level and ε is interpreted as model error which includes both model form error and solution approximation error. Statistics for the prior of ε could be obtained by subtracting the available model predictions (prior to updating) from the experimental data. It is acknowledged that experimental measurements can also contain errors. These could be addressed by adding terms to Equation 4.4 which become nodes in the Bayes network. This formulation can be directly implemented in WinBUGS and estimates of the posterior distribution of ε obtained.

One implementation issue is that to update some of the nodes in the Bayes network, evaluation of either the full finite element model or a surrogate model is needed. This issue needs to be addressed in the most efficient manner since multiple realizations of the model execution will be necessary. One leading candidate for this efficient calculation will be the use of surrogate models such as a Gaussian process (GP) model. Briefly, the

66

basic idea of the GP interpolation model is that the outputs, Y, are modeled as a Gaussian

process that is indexed by the inputs, X. A Gaussian process is simply a set of random

variables such that any finite subset has a multivariate Gaussian distribution. Once the

output is observed at m training points denoted as $x_1, \ldots, x_m$, the conditional distribution

of the process can be computed at any new point, x^*, which provides both an expected

value and variance of the surrogate model. A summary of GP modeling has been

presented in Chapter 2. Additional in-depth information can be obtained from McFarland

(2008) Rasmussen (1996), Martin and Simpson (2005), Mardia and Marshall (1984) and

Santner et al (2003).

4.3 Implementation and Results

A Bayes network has been developed for the demonstration problem as shown in

Figure 4.1. The baseline network includes all available data up to the second level and

quantifies the error related to model versus observed data that is not explained by

variability in the model parameters. The details of the implementation and some

preliminary results of the Bayes network approach to uncertainty quantification to the

problem shown in Figure 3.1 of Section 3.1 are described below.

A Markov chain Monte Carlo solution to the Bayes network shown in Figure 4.1

was found using the software WinBUGS (Spiegelhalter et al, 2003). Software to

implement the GP models was originally coded by McFarland (McFarland, 2008) and

Bichon (Bichon et al, 2008) in Matlab and Fortran. To make the GP evaluation software

available to WinBUGS, GP software was implemented as a function written in

Component Pascal (Oberon Microsystems, Inc., 2006) and compiled directly into the

WinBUGS software. This allows the software to make the necessary function evaluations which relate the model prediction to the calibration parameters and allow for the updating of the parameters within a Bayesian framework. The parameters to be updated are:

- Joints model parameters as described in Smallwood, et al (2000) and denoted as θ^j in Figure 4.1 are:

    - Linear stiffness, K_{lin}

    - Non linear stiffness, K_{non}

    - Degree of non linearity of energy dissipation vs force relationship, n_{pow}

- Foam model parameter denoted as θ^f in Figure 4.1 is modulus of elasticity, E.

- Error terms denoted as ε in Figure 4.1 are for levels 1 and 2, foam and joints (where $\varepsilon \sim N(0,\sigma)$).

Some selected results of the Bayes network implementation are shown below. For the estimation of probability density functions, 5,000 samples were used. These were in addition to 10,000 samples discarded as burn-in samples to allow the Markov chain to become stationary. Convergence of the Markov chain is assessed by considering the samples created after the burn-in. Figure 4.2 show the samples for the linear stiffness, K_{lin}, nonlinear stiffness, K_{non}, nonlinear exponent, npow and modulus of elasticity, E.
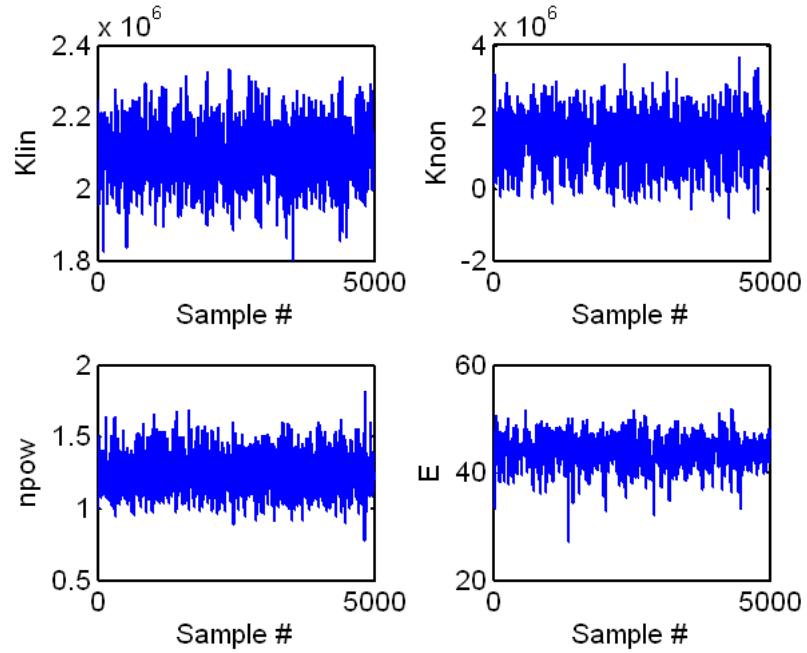
Figure 4.2. Samples of K_{lin}, K_{non}, npow and E used to demonstrate convergence of Markov chain to constant mean and variance

Convergence is checked by plotting in Figure 4.3, the moving average of the data from Figure 4.2 which shows converged mean values for K_{lin}, K_{non}, npow and E.
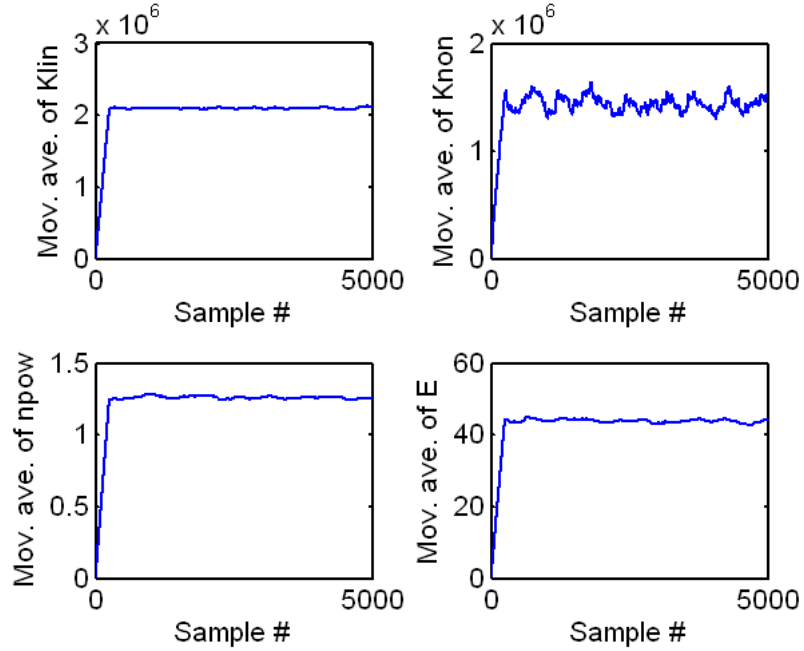
Figure 4.3. Moving average to assess Markov chain convergence

The first set of results show a comparison of the kernel density estimators (KDE) (Silverman, 1986) for the updated parameters of the joint model, K_{lin}, K_{non} and n_{pow} as they compare to the prior distributions. A KDE is an approximation to the probability density function (PDF) of source of values of s and it is computed from n data realizations, $s(.)_j \quad j = 1...n$. The form of the KDE used here is:

$$\hat{f}_S(\alpha) = \frac{1}{n} \sum_{j=1}^{n} \frac{1}{\sqrt{2\pi\varepsilon}} \exp\left[-\frac{1}{2\varepsilon^2}(\alpha - (s)_j)^2\right] \qquad -\infty < \alpha < \infty \qquad (4.5)$$

where ε is the "width" of a Gaussian kernel. Figure 4.4 through Figure 4.6 show these KDEs. Normal distributions were used as prior distributions for these parameters, with statistics (mean and standard deviation) obtained from the level 0 experiments. These are plotted as the solid lines in Figure 4.4 through Figure 4.6. These priors are updated with

the available energy dissipation data measured in levels 1 and 2. The effect of the available experimental data is reflected in each the posterior distribution relative to the prior. It is important to note that data from both levels 1 and 2 are simultaneously used to update these parameters and thus the effect of this data is reflected in the posterior of the parameters.
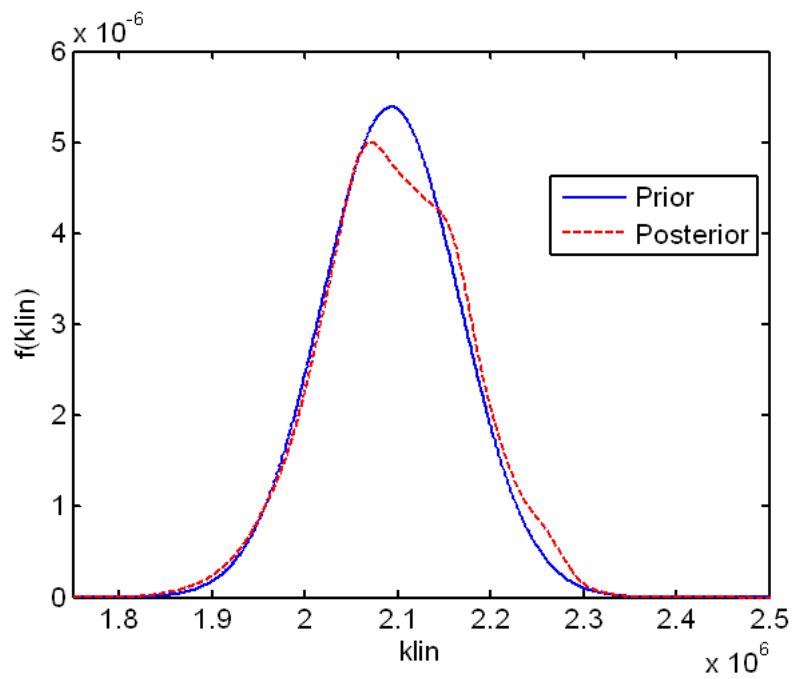
Figure 4.4. KDE of linear stiffness from the joint model
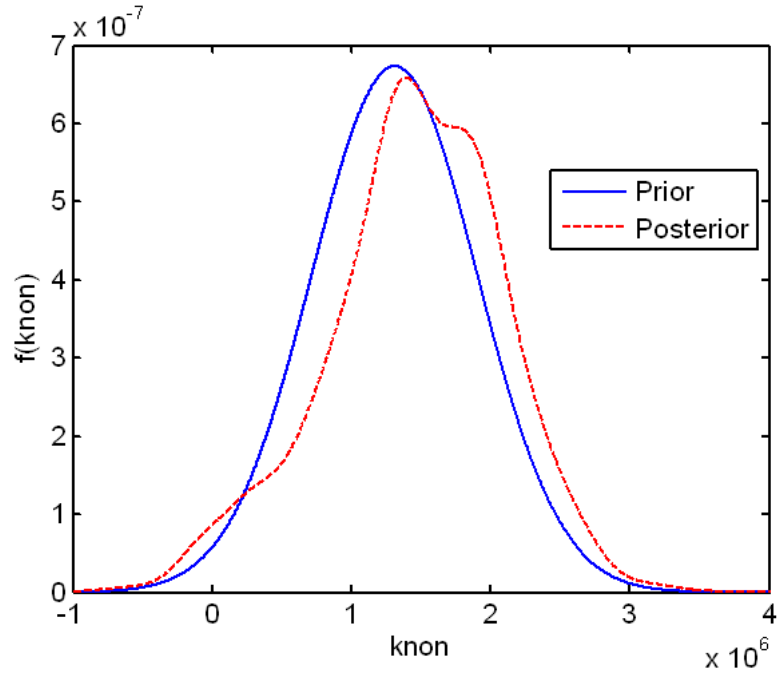
Figure 4.5. KDE on nonlinear stiffness from the joint model
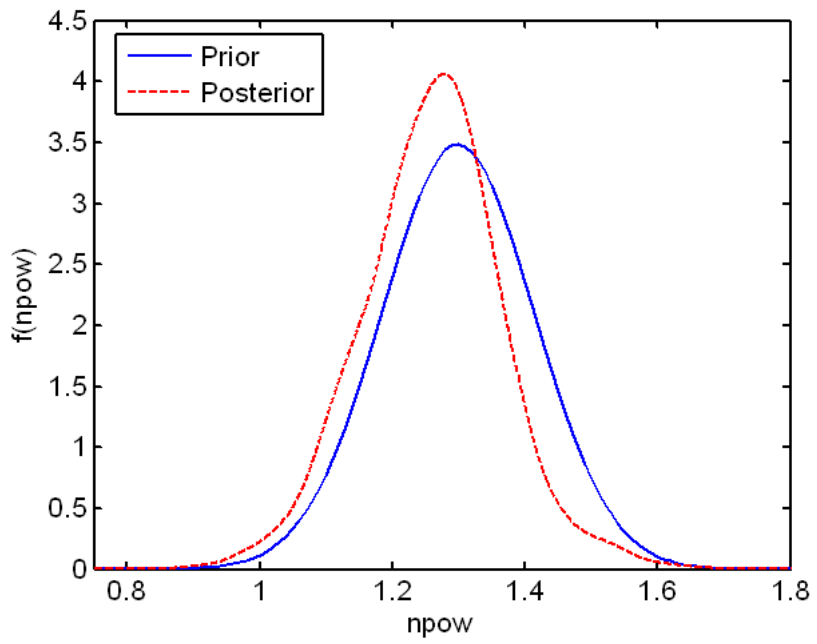
Figure 4.6. KDE of degree of nonlinearity from the joint model

Figure 4.7 shows the KDE for the modulus of elasticity which is the parameter of interest describing the foam. For the foam, the data used to update the parameter is the first natural frequency of the axial deformation mode measured at both levels 1 and 2. From the figure, it can be seen that after updating, the posterior distribution of the modulus of elasticity has a smaller variance when compared to the prior. This is a reflection of having more information and thus reducing the uncertainty about a given parameter. It also gives an indication of the possible range of values where this parameter might lie when both levels of complexity are included.
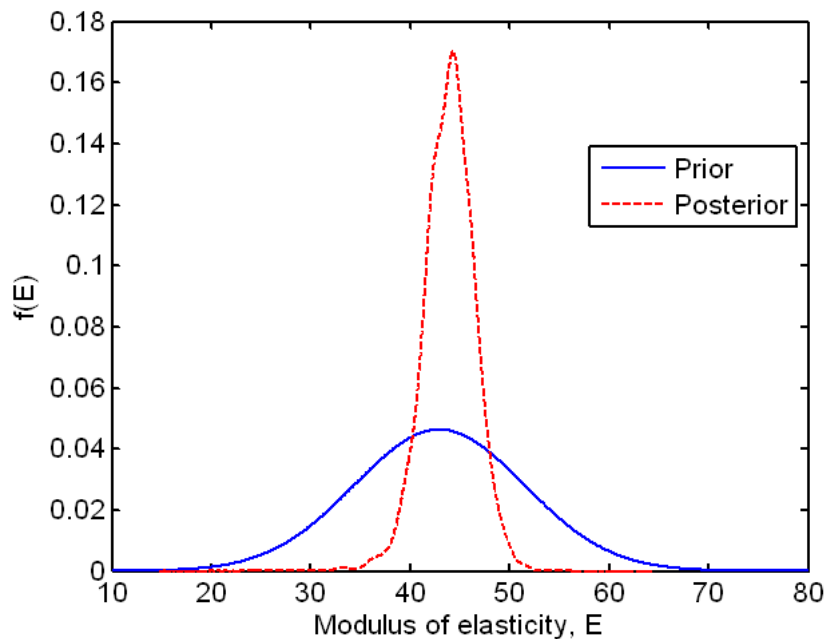
Figure 4.7. KDE of modulus of elasticity of foam

Figure 4.4 through Figure 4.7 showed the effects of the experimental data at several input parameter nodes in the Bayes network when updating the parameters. Next, we will see the effect on the predicted model responses relative to the actual data. It is important to point out that these plots are not intended to show a validation of the model. They are merely for comparison purposes. A more formal validation could be carried out, but it is not the focus of this study. For the joints, the quantity of interest is the energy dissipated at a particular force level. In this case, the force was chosen to be 300 lbs mainly because it is at the mid-point in the calibration data (level 0) which span 100 to 500 lbs. With this specification, Figure 4.8 shows the response quantity – energy dissipated per cycle of response at a mechanical joint force level of 300 lbs - for the level 1 hardware when computed using the prior distribution of the input parameters and after updating the parameters (shown as the posterior distributions). A comparison is also made to measured data. Figure 4.8 shows that the response of the model when evaluated using the updated parameters, is a good representation of the mean behavior of the hardware at level 1 when compared to the KDE of the experimental data. It can be observed that a large part of the PDF is located roughly around 1.3e-3 which is close to the mean of the test data and the variance of the prediction is smaller than the variance of the test data. It is also observed from the figure, the effect of the experimental data on the prior prediction. The net effect is a reduction in the variance of the posterior distribution of this parameter relative to the prior distribution.
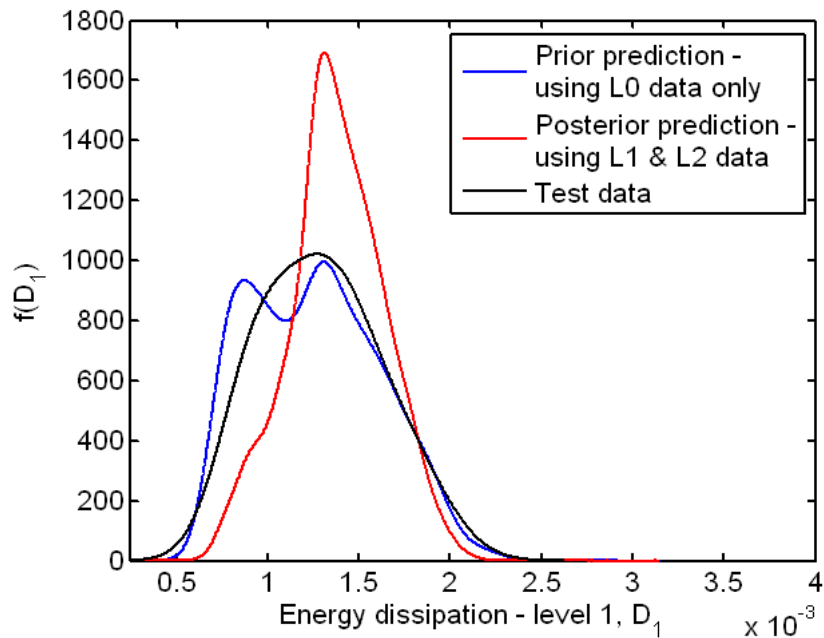
Figure 4.8. KDE of energy dissipation at level 1

Next consider the measure of response of the foam at level 2 which, in this case, is the natural frequency of the axial deformation mode of the hardware. The KDEs of the prior and posterior predictions and the posterior prediction including error are shown in Figure 4.9. Again, the prediction refers to the model response for the foam at level 2 using the updated parameter, (i.e. the posterior distribution of the modulus of elasticity) and the prior prediction is calculated from using the prior distributions of the parameters. Also shown are the actual test data used to update the prior distributions. These are shown as small circles plotted on the abscissa. In this case, the KDE of the predicted behavior encompasses 2 of the 3 data points. When the KDE of the prediction plus the error term is plotted, it now captures all available data. This indicates that there are other sources of uncertainty that come into play to explain the discrepancy between model

prediction and experimental data which are not fully explained by the model input parameters themselves. In this case the error term captures this discrepancy. This is an encouraging result since one of the hopes for this formulation was to properly apportion the uncertainty in model prediction to various sources. In other words, if a source of uncertainty is something other than parametric uncertainty, this should be captured as a separate error term, not rolled into the parameters themselves. In this research, no attempt is made (nor is it possible with the information available) to establish what the source of the error is. An understanding of the degree to which this error is present and whether it is a minor or major contributor to the overall uncertainty in the model prediction is desired. In this case, for level 2 foam, it shows some contribution. When observing Figure 4.9, it is clear that the KDE of the prediction plus the error spans the available data whereas the prediction by itself does not. Similarly to the energy dissipation at level 1, the prior prediction shows a larger variance when compared to the posterior distribution. This demonstrates one of the key features of a Bayes network and that is to incorporate data which has the effect of reducing the uncertainty of a quantity of interest, in this case, the natural frequency of the axial deformation.
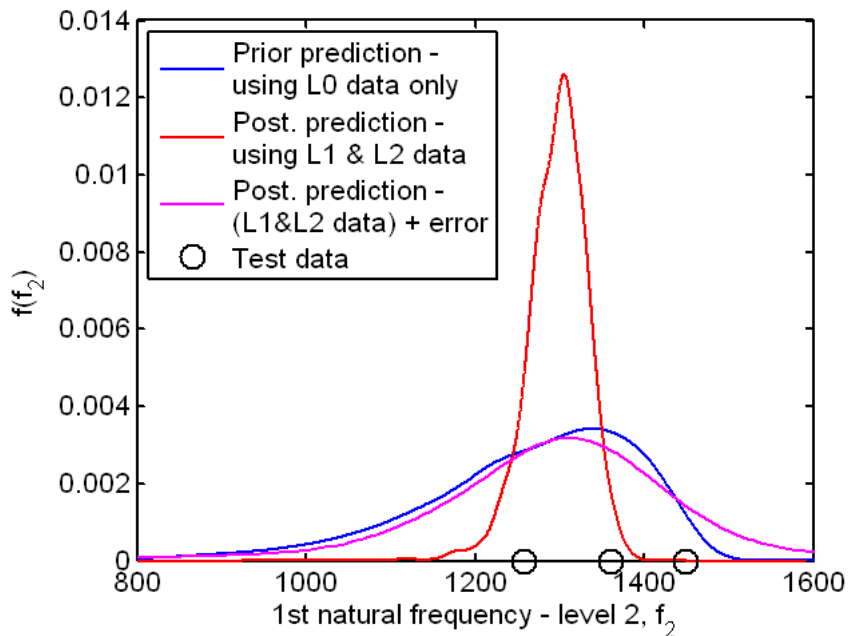
Figure 4.9. KDE of natural frequency of the axial mode at foam level 2 and test data
(shown as o)

Finally, the kernel density estimator of the (extrapolated) predictions at the system level is shown in Figure 4.10. The extrapolation is both in terms of the input to the system (excitation with a blast load) and in the system that is used (cone with encapsulated mass inside). Several plots are shown in this figure. The prior prediction refers to the system level response evaluated using the prior distributions of the input parameters. This KDE shows a slight bi-modality. The next 3 curves shown in the figure refer to the particular data used to perform the update. The "posterior prediction using L1 data only" means that only data from level 1 is used whereas "posterior prediction using L2 data only" refers to level 2 data being used for updating. The last curve represents the case where all the available data at all levels is used for updating the Bayes network. Among the observations that can be made from this figure, it is noted that the in general

all the posterior predictions show a decrease in variance relative to the prior which means that any available data reduces the uncertainty present in the system level response. In addition, and not surprisingly, when all the data is used, the variance of the posterior distribution of the system response is decreased the most relative to the other cases. From the figure, it can be observed the effect of the data from the various levels when it is included or not in the updating. It is seen from the figure that when removing the level 2 data, the variance of the response is similar to the case where all data is used for updating. In contrast, when level 1 data is not included, the variance increases relative to the case where all data is included. This is a type of sensitivity analysis which shows the effect of the data from the different levels on the variance of the system level response.
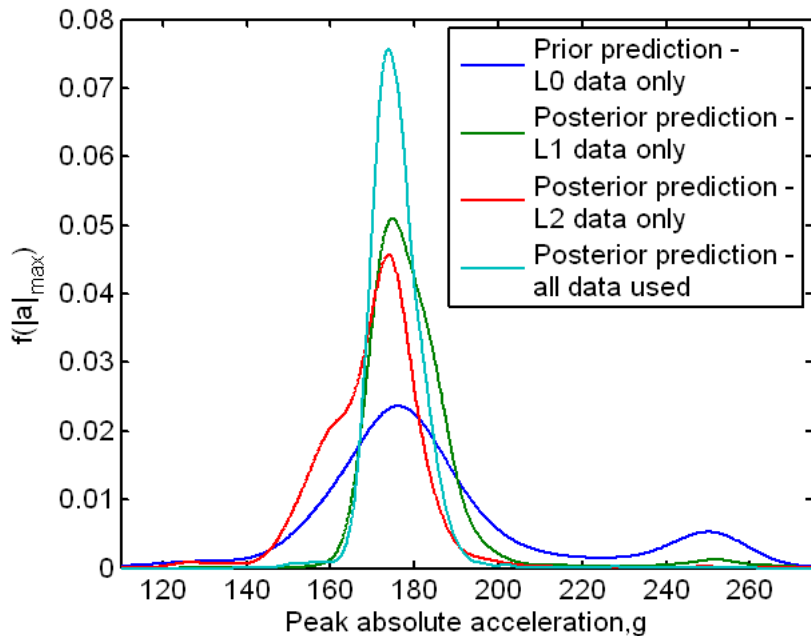
Figure 4.10. KDE of peak absolute acceleration at the system level

4.4 Model Errors

In this section, the errors between the model predictions and the experimental data at each level are analyzed. These errors are denoted as ε_1^j, ε_2^j, ε_1^f and ε_2^f corresponding to the joints level 1 and 2 and to the foam level 1 and 2 respectively, and are included in the Bayes network in Figure 4.1. The errors are assumed to follow a Normal distribution with zero mean and standard deviation σ. The standard deviation is updated with the observed data. This formulation is shown below:

$$\varepsilon \sim N(0, \sigma) \tag{4.6}$$

where the hyper-parameter σ is given a prior distribution of:

$$\sigma \sim \text{gamma}(0.0001, 0.0001) \tag{4.7}$$

and is subsequently updated in the Bayes network. In Equation 4.7, the first parameter of the gamma distribution refers to the shape of the distribution and the second parameter is the scale. All of the error terms are given the same prior for σ. The prior distribution for all the error terms is shown in Figure 4.11.
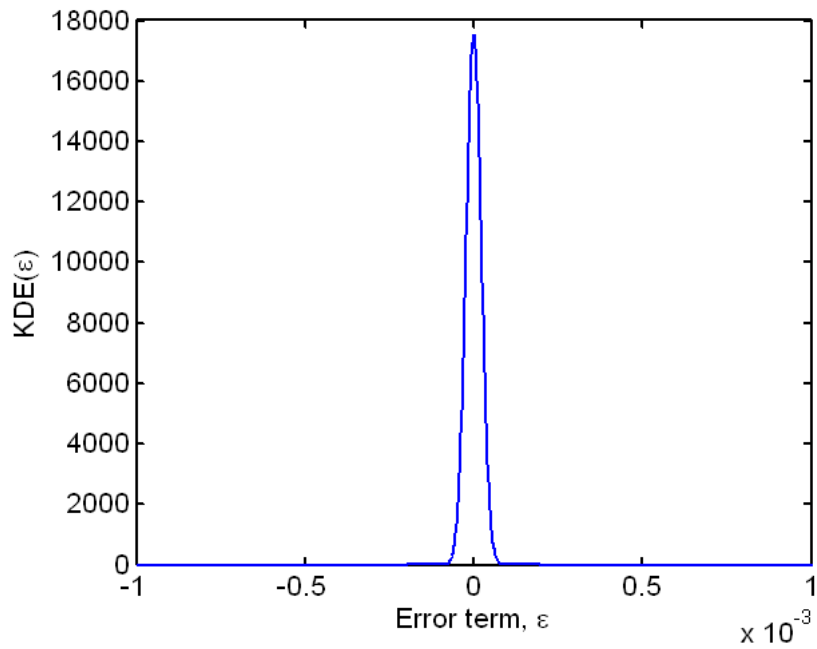
Figure 4.11. Prior distribution for all error terms in the Bayes network

After updating the Bayes network, the posterior distribution of all the error terms are calculated and plotted in Figure 4.12 through Figure 4.15. The priors are also plotted. Note that the priors have a very small variance compared to the posteriors and thus are plotted on a separate graph to help visualization.
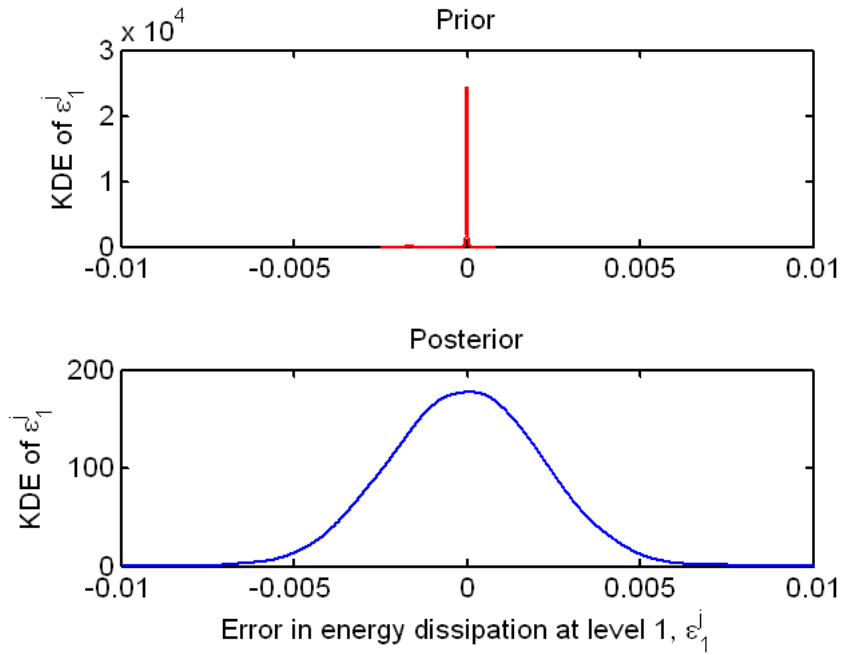
Figure 4.12. KDE of error term ε_1^j for joints at level 1
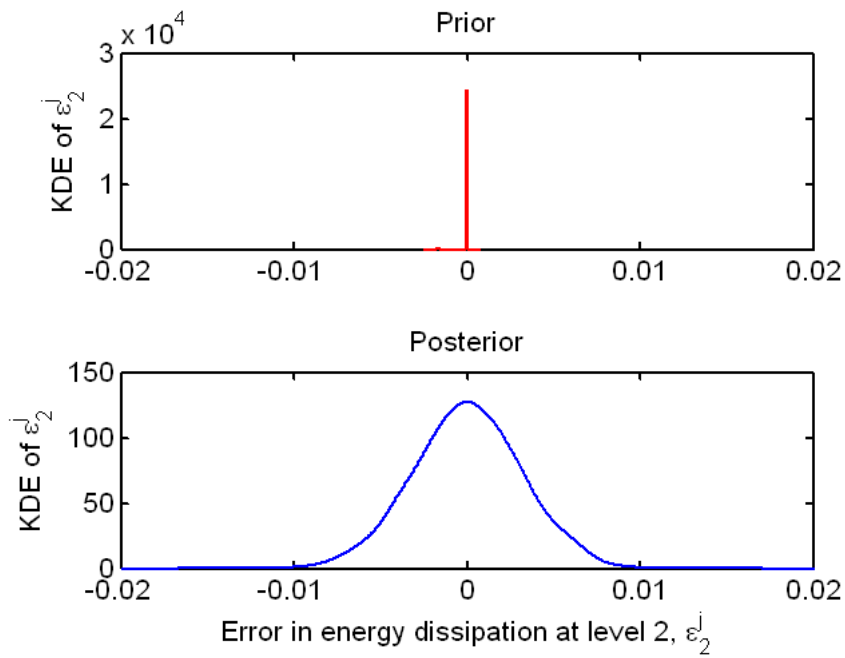
Figure 4.13. KDE of error term ε_2^j for joints at level 2

81
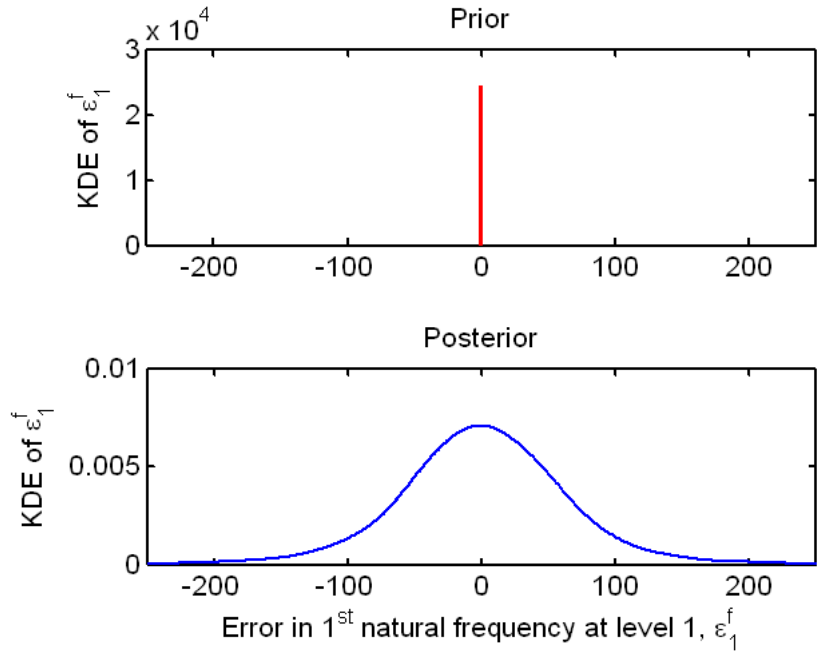
Figure 4.14. KDE of error term ε_1^f for foam at level 1
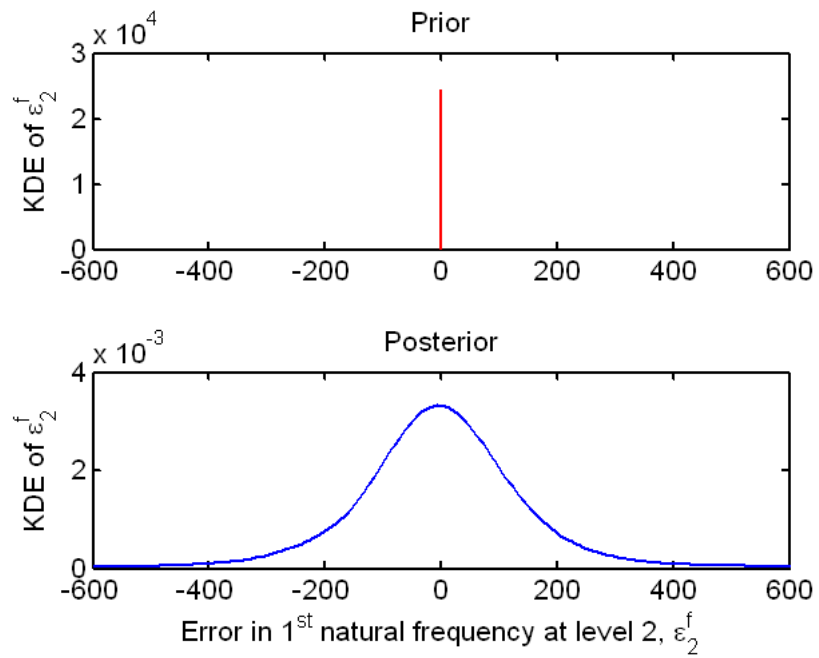
Figure 4.15. KDE of error term ε_2^f for foam at level 2

A general observation from these figures is that the model predictions are not perfect and some degree of error is present for all the measures of importance (i.e. energy dissipated and first axial frequency). Of importance in this study is the fact that this error term, which is a source of uncertainty, can be quantified and treated separately from the parametric uncertainty. Although not treated in this study, some possible sources of this error are:

1. measurement error,

2. model form error (arising from the choice of model selected to represent the physics of interest)

3. solution approximation error (e.g. mesh discretization)

An interesting topic for further research would be to examine the individual contribution of each error sources listed above to the overall error in the system level prediction. This will help identify possible areas of improvement and target those that have the greatest effect on the system level prediction error.

4.5 Conclusion

This chapter presents an approach to quantify and propagate uncertainty in a complex system model that is built in a hierarchical manner. This analysis incorporates sources of aleatoric uncertainty (input variables and model parameters) within the context of a multi-level, 2-component problem using a Bayes network as the main analysis framework. The Bayes network is chosen because it conveniently allows the modeling of

both causal relationships and statistical dependencies between sets of data within a probability framework, based on conditional probabilities between the various nodes of the network. This framework also allows any available experimental data at any level to be incorporated into the analysis and calculates the posterior probabilities of all nodes in the network.

A key issue that is addressed is the actual implementation of this methodology to perform fast and efficient computation within a Bayesian framework. This is done by using surrogate models (specifically Gaussian process models) in lieu of the finite element analysis which can be expensive to run. Results show that using a Bayes network approach is a reasonable way to model a multi-level problem and available experimental data can be easily incorporated into the analysis. Error terms were included in the prediction at each level to account for model errors in addition to parametric uncertainty alone. These error terms were quantified but their sources were not determined at this stage.

The next step is to incorporate epistemic uncertainty into the analysis via the model parameters. This uncertainty can arise from lack of knowledge about a parameter of interest and will need to be treated probabilistically in order to be included in the Bayes network. Following this, the contribution of different sources of uncertainty to the overall uncertainty in the system level prediction could be investigated.

CHAPTER V

INCORPORATION OF BOTH EPISTEMIC AND ALEATORIC UNCERTAINTY IN A HIERARCHICAL SYSTEM MODEL

5.1 Overview

This chapter presents an extension of the methodology described in Chapter 4 that adds epistemic uncertainty into the Bayes network. As described in Section 2.4, sources of epistemic uncertainty arise from lack of data and knowledge. In the context of this research, there can be limited knowledge of a material or component behavior, the coupling of two or more component models, the extrapolation of a system model to an application space and generally, any condition that the modelers did not anticipate occurring and thus, is not included in the model's expected use. In this chapter, two main sources of epistemic uncertainty are considered: data uncertainty and model error. Data uncertainty is treated in a probabilistic manner via a family of flexible distributions known as the Johnson distribution. As in Chapter 4, model error is quantified and discussed but its sources are not investigated.

5.2 Background

Data uncertainty is introduced with respect to the model parameters and arises from the fact that some materials used in engineering systems may not be fully characterized. Usually their parametric description is given in terms of intervals defined by subject matter experts and/or by very limited information but not by full probability distributions.

Formally, Ferson (2004 and 2007) and Osegueda et al (2002) list eight sources from which information is best represented by intervals, including plus-or-minus reports, significant digits, intermittent measurement, non-detects, missing data and gross ignorance. Intervals are obtained by either examination of the limited information available (enough to establish bounding information but not a full probabilistic representation) or by eliciting information from subject matter experts. For the case of expert opinion, one should also consider the relative weight that each expert's opinion carries and it should be incorporated into the uncertainty analysis. This assessment of weight is, for the most part, a subjective endeavor and a future work topic. In this research, the parameter describing the foam behavior will be treated as a source of epistemic uncertainty and is assumed to be given as intervals by subject-matter experts. The approach taken in this research uses a transformation of the data uncertainty into a probabilistic description and it is detailed in McDonald et al (2009a, b and c). This approach starts by calculating the mean and estimates of the next three central moments of the bounds of intervals representing the epistemic uncertainty. These moments are used to estimate the parameters of the Johnson family of distribution by one of several methods.

Another source of uncertainty is the choice of method for the treatment of epistemic uncertainty in the probabilistic calculations. In this study, this type of uncertainty arises from the various methodologies available to fit parameters of a Johnson distribution to the given data. This effect has not been considered before, and it is explored in this work by making comparisons of the results using the decision-making metric of interest when the various methods are implemented.

5.2.1   Estimating Bounds on Moments

The estimation of statistics or moments of interval data has been the subject of several recent papers and formulas and algorithms necessary to compute basic statistics for interval data have been presented. (e.g., Ferson et al, 2007; Kreinovich et al, 2004; Gioia and Lauro 2005; Xiang et al, 2006). In this research, the methodology developed in McDonald et al (2009a, b and c) is used to estimate the lower and upper bounds of moments from the interval data. The formulas to calculate the bounds on the first four moments given interval data are shown in Table 5.1.

Table 5.1. Formulas to calculate first four moments of interval data

Moment	Formula for lower and upper bound
First Moment	$$\left[\underline{M},\overline{M}\right]=\left[\frac{1}{N}\sum_{i}^{N}a_i,\frac{1}{N}\sum_{i}^{N}\overline{b}_i\right]$$
Second Moment	$$\min/\max f(x)=V=\frac{1}{N}\sum_{i=1}^{N}x_i^2-\frac{1}{N^2}\left(\sum_{i=1}^{N}x_i\right)^2$$ $$s.t.\quad x_i\geq a_i$$ $$x_i\leq b_i$$
Third Moment	$$\min/\max f(x)=M_3=\frac{1}{N}\sum_{i=1}^{N}(x_i-\frac{1}{N}\sum_{j=1}^{N}x_j)^3$$ $$s.t.\quad x_i\geq a_i$$ $$x_i\leq b_i$$
Fourth Moment	$$\min/\max f(x)=M_4=\frac{1}{N}\sum_{i=1}^{N}(x_i-\frac{1}{N}\sum_{j=1}^{N}x_j)^4$$ $$s.t.\quad x_i\geq a_i$$ $$x_i\leq b_i$$

5.2.2 Johnson Family of Distributions

As a basis to construct a probabilistic representation, the Johnson family of distribution (Johnson, 1949) is used in this research. The reason for choosing Johnson's family of distributions is that it is a flexible set of probability distributions capable of representing a wide array of conventional probability distributions (McDonald et al (2009a, b and c)). If X is a continuous random variable with distribution function $F(x) = P(X \leq x)$, Johnson (1949) proposed four normalizing translations of the general form:

$$Z = \gamma + \delta \cdot g\left(\frac{X - \xi}{\lambda}\right) \tag{5.1}$$

where Z is a standard normal random variable, γ and δ are shape parameters, λ is a scale parameter, ξ is a location parameter and g(.) is a function that defines the four distribution families as:

$$g(y) = \begin{cases} \ln(y) & \text{for } S_L \text{ (lognormal) family} \\ \ln\left(y + \sqrt{y^2 + 1}\right) & \text{for } S_U \text{ (unbounded) family} \\ \ln(y/(1-y)) & \text{for } S_B \text{ (bounded) family} \\ y & \text{for } S_N \text{ (normal) family} \end{cases}$$

$$\tag{5.2}$$

If sample data is available, choosing which family of distributions to use is accomplished by the following procedure:

1. Estimate the first four central moments, m_1, m_2, m_3, and m_4.of the sample data, X as: (DeBrota et al, 1998)

88

$$m_1 \equiv E(X) \text{ and } m_k \equiv E(X - m_1)^k \quad k = 2, 3, 4. \qquad \text{(5.3 and 5.4)}$$

where E(.) is the expected value of the quantity inside the parenthesis.

2. Calculate the skewness and kurtosis:

$$\beta_1 \equiv m_3^2 / m_2^3 \text{ and } \beta_2 \equiv m_4 / m_2^2. \qquad \text{(5.5 and 5.6)}$$

3. Use the identification chart in Figure 5.1 to determine the appropriate distribution family to use.
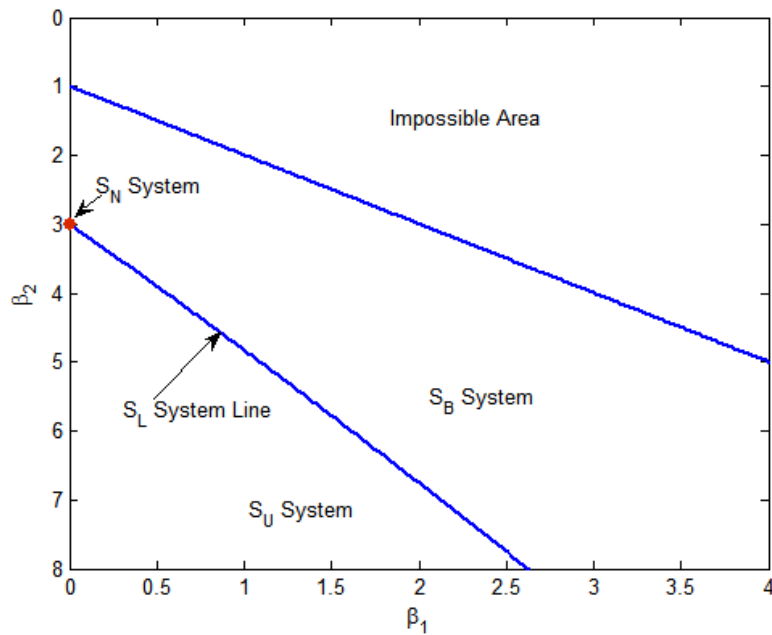
Figure 5.1. Johnson distribution family identification chart

4. To fit the distribution parameter λ, ξ, γ and δ for one of the family of distributions shown in Equation 5.1, the following methods could be used

- method of matching moments (using the first four moments of the data),

- percentile matching (where a desired value is specified at a given percentile point(s))

- least squares estimation and

- minimization of the error norm of Johnson distribution when compared with empirical cumulative distribution function (CDF).

In this study, modifications of the first two methods will be used and are implemented by either solving a least squares or an optimization problem.

In the chart, S_U represents an unbounded distribution (support is $\pm \infty$) and S_B is a bounded distribution. The bounded family of distributions will be used for the problem examined in this research for the simple reason that the physics of the problem establishes natural bounds on the variable being modeled. That is the modulus of elasticity has to be greater than zero and it has a physical upper limit.

To use the Johnson family of distributions when interval data is available and when a moment based approach is desired, the process outlined above can still be applied except that instead of calculating moments from sample data, the moments would come from sampling within the bounds on the moments using the formulas presented in Table 5.1. Various methodologies are considered and implemented in this work to find the family of distributions that would model the epistemic uncertainty using bounds on moments. In addition, other methods are also investigated that do not rely on bounds on moments or are based on the Johnson distribution. The methods used in this research can be divided in two categories: methods that aggregate the interval data and those that treat

the interval data individually. The aggregation methods seek to combine the given interval data using the bounds on the moments obtained from the intervals and/or selected percentile values. The latter methods treat each expert as a separate entity and do not combine the interval data. The methods are briefly summarized below and more detail is given in the following sub-sections with an application of each method to the demonstration problem.

Methods for aggregating interval data

1. Method of matching and bounding moment – This method seeks to construct a family of distributions, using a parametric definition based on a Johnson distribution for which the first and second moments of a proposed distribution will match exactly a given set of sampled moments from the expert given bounds and the third and fourth moments will lie between a given set of bounds. This is a slight variation of what is proposed in DeBrota et al (1998).

2. Method of percentile matching and mean bounding– This method also constructs a family of Johnson distribution which are characterized by 2 selected percentile values defined by the empirical CDF obtained from the expert opinion being matched exactly and a selected mean falling between 2 given bounds. This is also a variation of what is proposed in DeBrota et al (1998).

Methods for individual treatment of interval data

1. Parametric distribution functions from each expert interval – For this method, a family of distributions is obtained by randomly selecting plausible values of the

distribution parameters and using the individual bounds of each expert to define a distribution.

2. Uniform distribution from each expert interval – This is similar to #3 above but it only defines a uniform distribution for each of the intervals.

3. One uniform distribution covering all the experts' range – This is the most un-informed choice of prior. It uses the minimum and maximum values considering all the expert given bounds to define the limits of a single uniform distribution.

One important issue to keep in mind when evaluating the various techniques to describe the epistemic variable is that each of these methods only produces a prior distribution for a given parameter. This prior distribution will be updated with any available data through the use of the Bayesian network construct which was described in Chapter 4. When sufficient data is incorporated into the analysis during the Bayesian updating process, the posterior probability distribution tends to have diminishing dependence on the prior distribution, regardless of the form of the prior. It is important to note that the priors should be defined over a range that includes values of the likelihood function (i.e. where data is available); otherwise, the resulting posterior distribution will resemble a delta function.

5.3 Treatment of interval data: Methodology and Implementation

This section considers various methodologies to handle interval data in the parameters of a model. Since the objective is to treat the entire problem in a probabilistic way, the objective will be use the available interval information and assign a reasonable

probability distribution that both incorporates the given information (i.e. bounds) and also

does not add too much subjective information. For this research the following intervals

on the modulus of elasticity, E given by six subject matter experts are given as:

Table 5.2. Lower and upper limits for modulus of elasticity, E, from six experts

Expert	Lower Limit (ksi)	Upper Limit (ksi)
1	32	60
2	35	68
3	40	72
4	42	78
5	48	82
6	50	94

In this chapter, it is assumed that the foam characterization data (collected at level 0) is

not available and it is replaced by the experts' interval data shown in the table above. For

comparison purposes, the values of E used in Chapter 4 range between 20 ksi and 70 ksi

whereas the experts' range from 32 ksi to 94 ksi. The range of values given by the experts

are their best estimate of the parameter E at the system level and does not necessarily

need to coincide with the available experimental data.

For the joints, the level 0 data is still available and thus can be used to construct

probability density functions of the parameters of the joint model. This was shown in

Chapter 4.

With the intervals given in Table 5.2 and using the formulas in Table 5.1, the bounds on the first four moments are estimated and are shown in Table 5.3.

Table 5.3. Bounds on the first four moments from interval data

Moment	Lower Bound	Upper Bound
1	41.2	75.7
2	0	6288
3	-7800.4	14392.5
4	0	592721.8

5.3.1   Method of matching and bounding moments

This method uses the bounds on the moments calculated from the experts' intervals using the formulation described in the previous section and fits a bounded Johnson distribution from which samples can be drawn. A variation of the method of matching moments proposed by DeBrotra et al (1998) is used and to fit a bounded Johnson distribution to these moments, the implementation developed by Venkataraman and Wilson (1997) is adopted. In this methodology, the distribution is assumed to have the same moments about the origin as the observed data. The procedure to create realization of a bounded Johnson distribution based on the interval data shown in Table 5.2 is described below:

1. Calculate bounds on 1st moment via averaging of given bounds and on 2nd, 3rd and 4th moments via optimization, using the formulas shown in Table 5.1. These bounds on moments are shown in Table 5.3.

2. To generate the parameters of a bounded Johnson distribution

   a. Sample moments independently using the bounds on moments obtained in step 1 as the lower and upper bound of a uniform distribution. This is shown in Figure 5.2.
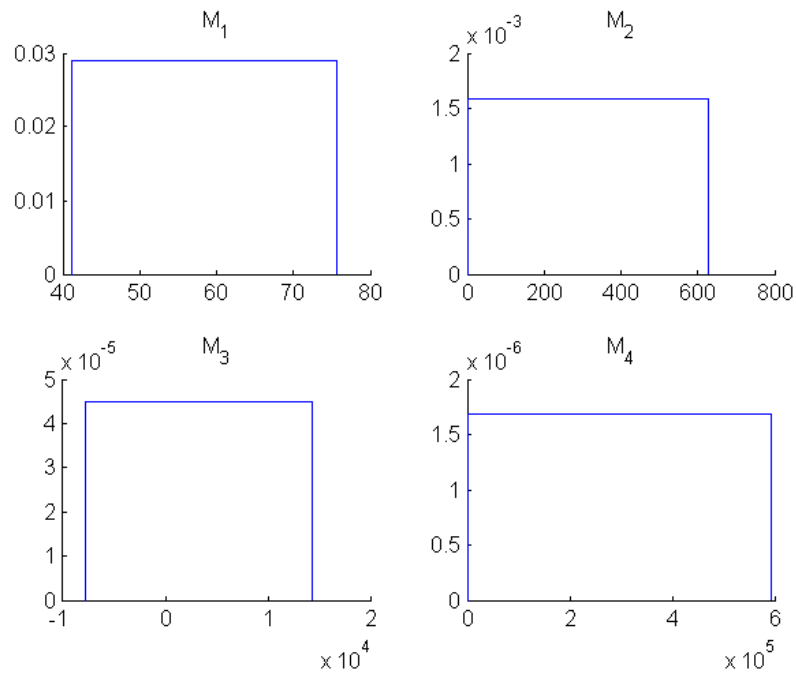
Figure 5.2. Moments are sampled from a uniform distribution with limits given by estimated bounds on moments

   b. Form a collection of n four-tuple moments sampled at random from the uniform distributions shown in Figure 5.2.

$$\mathbf{m} = \begin{bmatrix} m_1(1) & m_2(1) & m_3(1) & m_4(1) \\ \vdots & \vdots & \vdots & \vdots \\ m_1(n) & m_2(n) & m_3(n) & m_4(n) \end{bmatrix} \tag{5.7}$$

c. For each row in **m** in Equation 5.7, use Equations 5.5 and 5.6, calculate β_1 and β_2 and identify the region in Figure 5.1 where these fall. Keep only those combinations of β_1 and β_2 which fall in the bounded region.
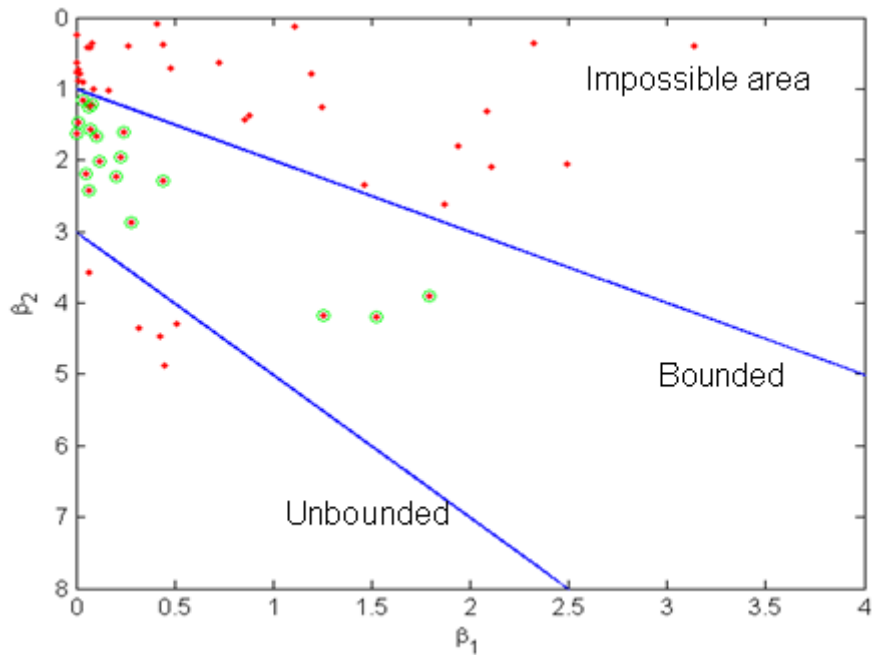
Figure 5.3. System identification plot showing combinations of β_1 and β_2 which yield a bounded system (shown as green circles)

d. Now, from the form of the bounded Johnson distribution

$$X = \xi + \lambda \cdot f^{-1}\left(\frac{Z-\gamma}{\delta}\right) \qquad \text{where} \quad \xi < X < \xi + \lambda$$

and

$$f^{-1}(z) = \frac{1}{\left(1+e^{-z}\right)}$$

(5.8)

and using the given interval data, the parameters λ and ξ can be established.

In this case they are set as:

ξ is min(lower bound) = 32
λ = max(upper bound) - ξ
λ = 94 - 32
λ = 62

e. Now use the method of matching moments to estimate remaining parameters, δ and γ. This is done using an optimization formulation as shown below:

$$\min \quad \text{f}(\delta,\gamma) = \sum_{i=1}^{2}\left(\frac{m_i - \hat{m}_i}{m_i}\right)^2$$

where $\qquad \hat{m}_i = E(x - \mu_x)^i$

$$X = \xi + \lambda \cdot \text{f}^{-1}\left(\frac{Z-\gamma}{\delta}\right), \qquad \xi < X < \xi + \lambda$$

and $\qquad \text{f}^{-1}(z) = \frac{1}{\left(1+e^{-z}\right)}$

(5.9)

s.t.
$$M_{3L} \leq \hat{m}_3 \leq M_{3U}$$
$$M_{4L} \leq \hat{m}_4 \leq M_{4U}$$

where m_j, $j = 1...4$ are the rows of **m** in Equation 5.7, $\hat{m}_j$, $j = 1...4$ are moments from sampled data taken from a bounded Johnson distribution with parameters λ, ξ, γ and δ, $\text{f}(\delta,\gamma)$ is the objective function and it is a sum of normalized squared errors between the 1st and 2nd moments,

97

m_l and $\hat{m}_l$, $l = 1, 2$ and M_{3L}, M_{3U}, M_{4L} and M_{4U} are the lower and upper

bounds on the 3rd and 4th moment respectively.

3. The above procedure creates n bounded Johnson distributions.

The formulation in Eq. 5.9 requires that a candidate bounded Johnson distribution has 1st

and 2nd moments that match those sampled from within the bounds shown in Table 5.3;

additionally, the 3rd and 4th moments of this candidate distribution fall within their

respective bounds (also shown in Table 5.3). This formulation is implemented in Matlab

using fmincon to solve the optimization problem and the parameters of bounded Johnson

distributions are obtained. Samples from these distributions are obtained and their PDFs

are shown in Figure 5.4.
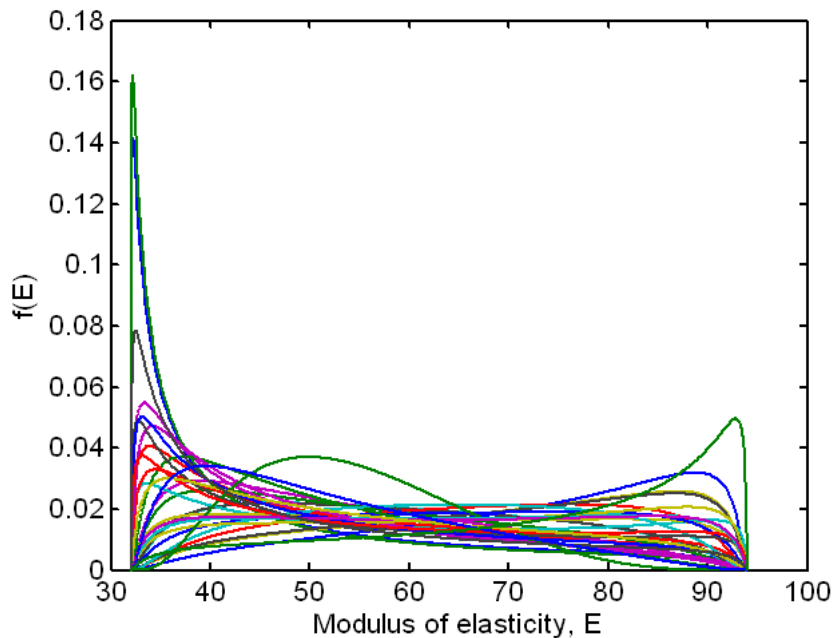
Figure 5.4. PDFs of bounded Johnson distributed foam modulus of elasticity with
parameters obtained using moment matching technique

To confirm that the simulated PDFs satisfy the required constraints, i.e. 1^{st} and 2^{nd} moments are matched and 3^{rd} and 4^{th} moments fall within the estimated bounds (as shown in Table 5.3), the moments from the PDFs shown in Figure 5.4 are calculated (via sampling from these PDFs) and plotted below. First, we look at the first moment. There is a 1% or less error between the calculated and target moments.
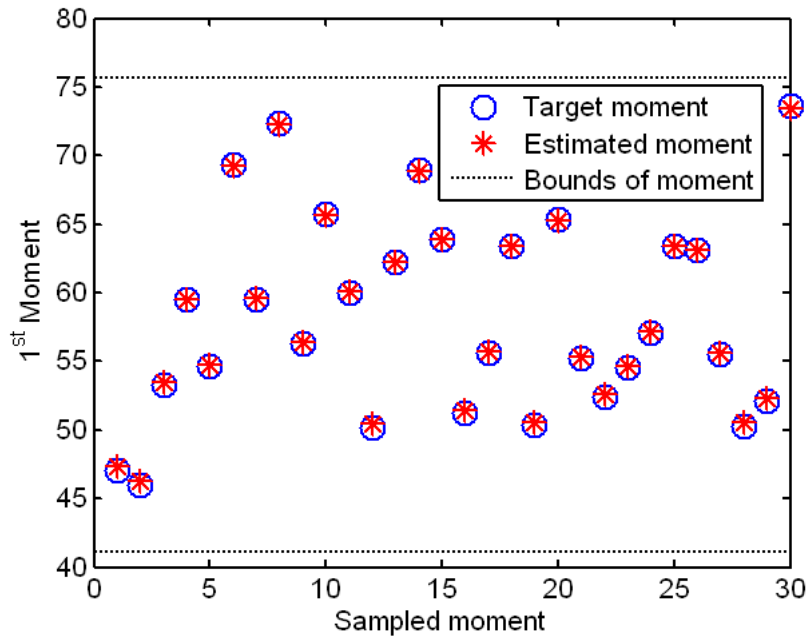
Figure 5.5. Comparison of 1^{st} moment sampled from Eq. 5.7 and those calculated from PDFs shown in Figure 5.4

Figure 5.6 shows the comparison between the calculated and target 2^{nd} moments for the PDFs. In this case, the RMSE is around 12.9 and the estimated quantities show a slight bias. This is a function of the tolerance of the optimization algorithm, in this case the

active set algorithm implemented in fmincon in Matlab, as it seeks to simultaneously satisfy the constraints that both the 1st and 2nd moments be matched. When this error is compared to the magnitude of the mean of the target moments, it is on the order of 5%, which is relatively low and, therefore, deemed acceptable for the purpose of this research.
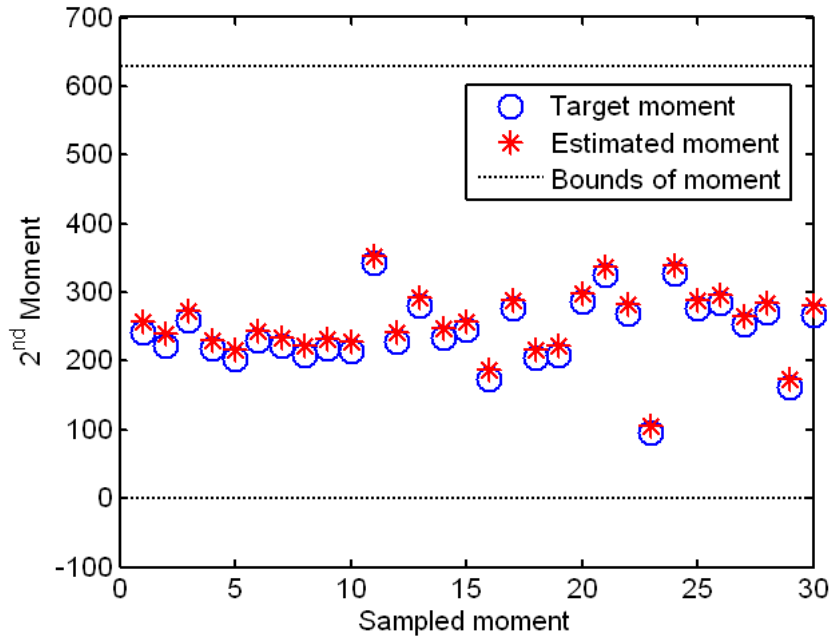
Figure 5.6. Comparison of 2nd moments sampled from Eq. 5.7 and those calculated from PDFs shown in Figure 5.4

The next quantities to be compared are the 3rd and 4th moments which are only required to fall within the calculated bounds. These moments were not required to be matched exactly (or at least to within some small error) because it was deemed that the main features of the distribution that need to be matched closely were the mean and the variance. Figure 5.7 and Figure 5.8 show the 3rd and 4th moments calculated from the

PDFs shown in Figure 5.4 relative to their bounds. As can be observed in these figures, the PDFs shown in Figure 5.4 satisfy the 3rd and 4th moment requirement and thus the PDFs are considered a good representation of the expert specified intervals shown in Table 5.2 and estimated via moment matching and bounding. These PDFs can be used as prior distributions for updating within the Bayes network.
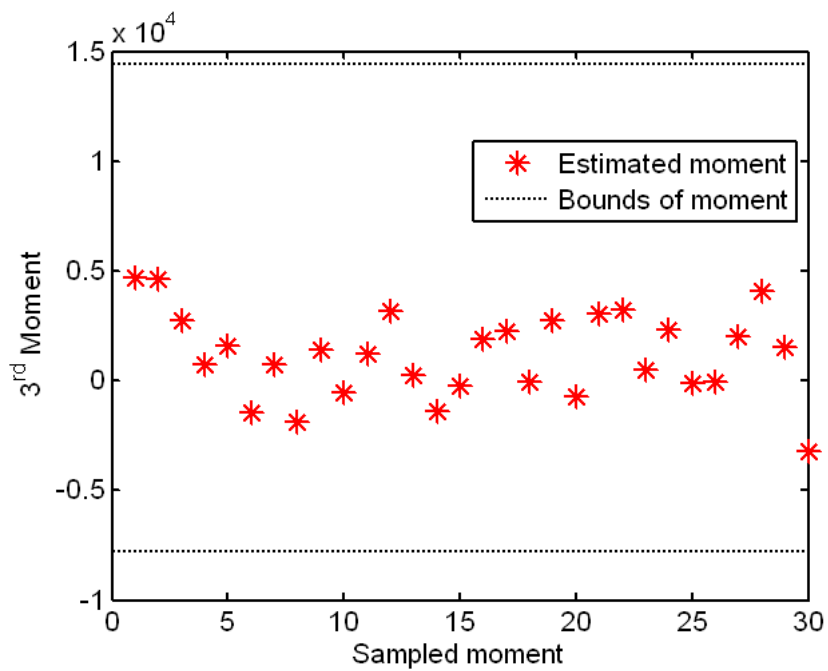
Figure 5.7. Comparison of 3rd moments calculated from PDFs shown in Figure 5.4 and its bounds
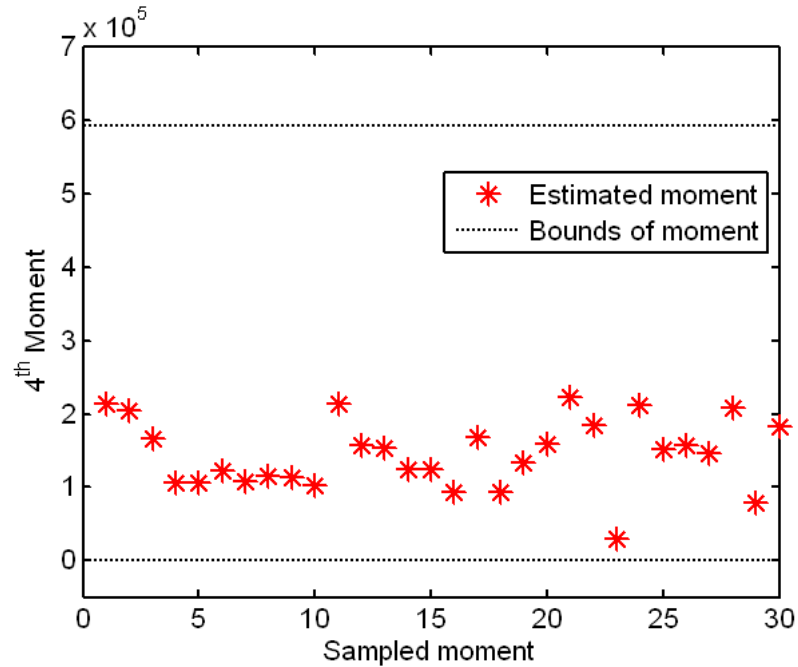
Figure 5.8. Comparison of 4^{th} moments calculated from PDFs shown in Figure 5.4 and its bounds

While observing Figure 5.6 through Figure 5.8, it is noted that the 2^{nd} to 4^{th} moments plotted on these figures are concentrated in a region which in general fall away from one or both bounds. Even though the sampled moments come from a uniform distribution which cover the space between the bounds of the moments, the requirements is that the combination of the moments fall in the bounded region of the identification chart shown in Figure 5.1 and these moments are used to fit a bounded Johnson distribution to. In addition, the optimization procedure used to estimate the parameters of the bounded Johnson distribution (Eq. 5.8) may not converge for all the candidate moments. Thus from the original space of candidate moments selected, only a subset fulfill all the requirements needed (i.e. β_1 and β_2 fall in the bounded region and

convergence of the optimization problem for estimating the parameters of the bounded Johnson distribution).

Once these distributions are created and it is verified that they satisfy the given requirements in terms of moments being matched and bounded, they now become the prior distributions to be used in the Bayes network shown in Figure 4.1. These priors are shown again below to facilitate the presentation of results.
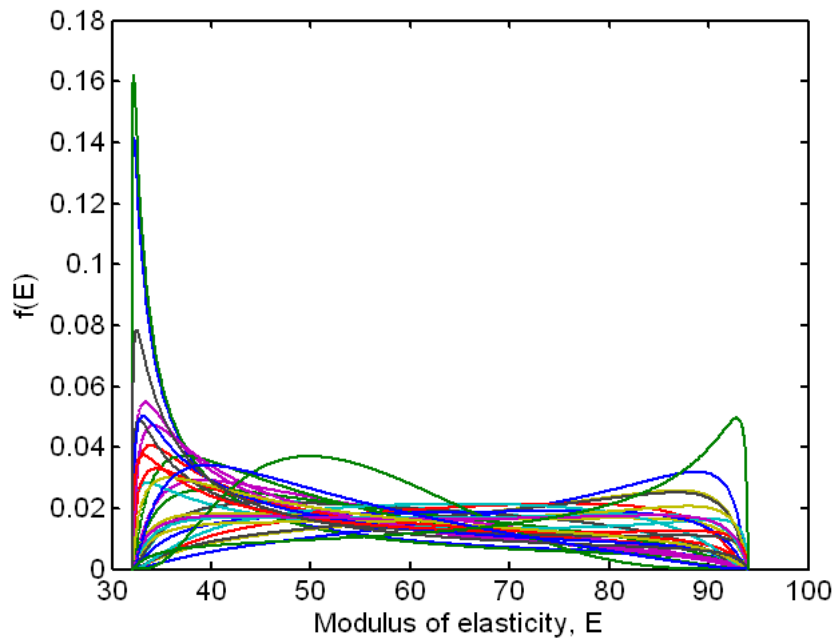
Figure 5.9. PDFs of prior distributions of foam modulus of elasticity, *E*

These priors are then updated using the level 1 and 2 data described in Chapter 3 and a posterior distribution of *E* can be obtained. These are shown in Figure 5.10.
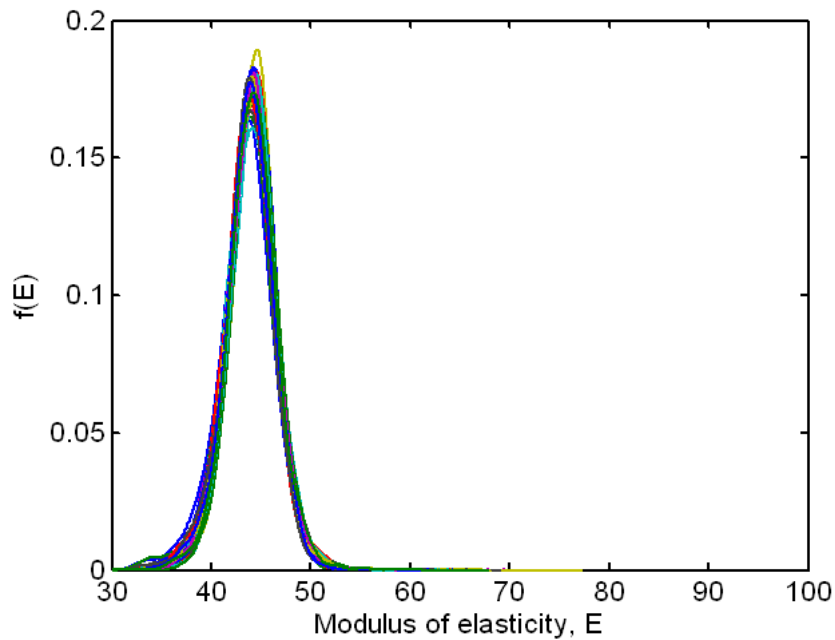
Figure 5.10. KDEs of posterior distributions of foam modulus of elasticity, *E*

The effect of data used for calibration is noted when comparing Figure 5.9 with Figure 5.10. The large variances in the PDFs of the priors are significantly reduced when they are updated to form posteriors; it is evident in the smaller variances of the posterior distributions. The posteriors converge to a collection of very similar KDEs.

There is a similar updating of the parameters for the joint model. These are treated as a source of aleatoric uncertainty where the statistics of the priors are obtained from a large dataset at level 0 and the priors are assumed to follow a Normal distribution. In the case of the joint parameters, there are significantly more data available for updating than for the foam. The posterior probabilities of the joint parameters in general look very similar to those shown in the results section of Chapter 4 so they would not be presented

here. The focus of this section is the source of epistemic uncertainty, the foam's modulus of elasticity.

The final step in this process is to obtain a forward prediction of the system level response, the peak acceleration of the encapsulated mass. These are shown in Figure 5.11. There is one peak acceleration response KDE for each prior PDF on the foam modulus of elasticity, E. As it is evident from Figure 5.11, the distributions of the peak acceleration at the system level are very tightly grouped, as a matter of fact they almost seem to converge to the same KDE. This is a reflection the posterior distribution of the foam modulus of elasticity which also shows a tight grouping and the fact that the response of the encapsulated components is very sensitive to this foam modulus of elasticity.
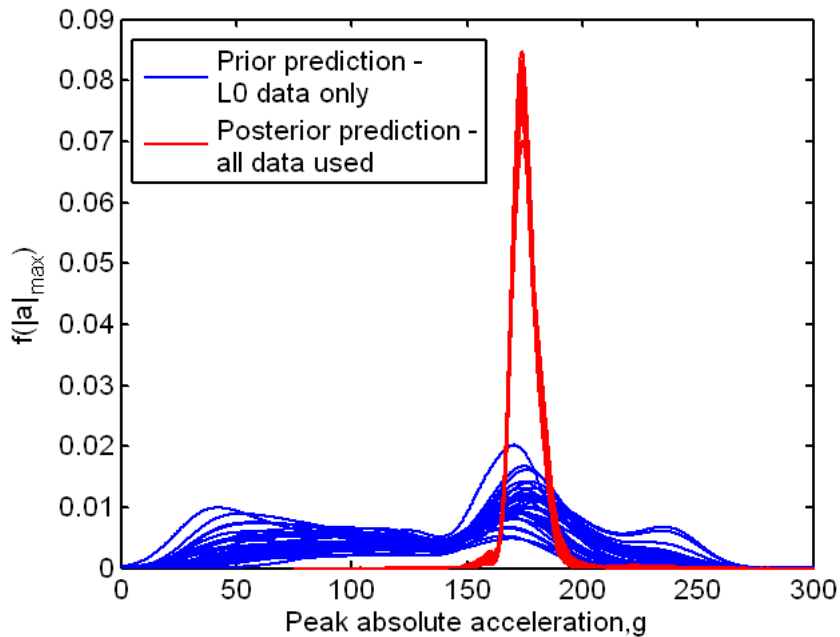
Figure 5.11. KDEs of peak accelerations at the system level

The final plot shows a comparison when only sources of uncertainty are considered aleatoric only versus both aleatoric and epistemic are included. This is shown in Figure 5.12. The results for this technique are very similar to those obtained when aleatoric uncertainty is only considered. This could indicate that the priors obtained with the method of moment matching share similar characteristics (i.e. 1st four moments) of those of the modulus of elasticity when it is treated as an aleatoric variable. This leads to a posterior prediction at the system level that are very similar to each other.
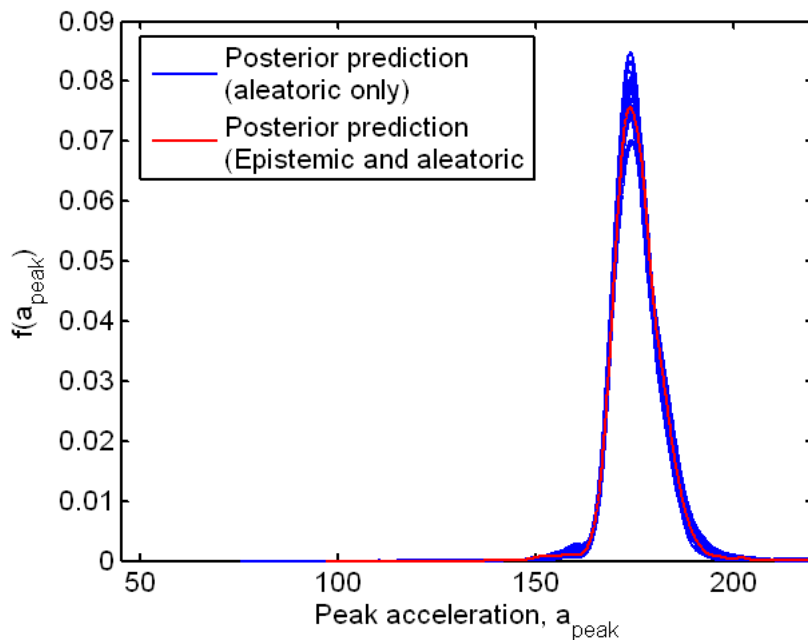
Figure 5.12. Comparison of system level predictions when aleatoric only and both aleatoric and epistemic uncertainties are included

5.3.2 Method of percentile matching and mean bounding

This method starts with the bounding information given by subject experts and computes the empirical cumulative distribution function (ECDF) of the experts' lower limit specifications and the ECDF of the experts' upper limit specifications. The ECDF is a cumulative distribution that concentrates probability $1/n$ at each of the n numbers in a sample (Cox and Oakes, 1984 and The Mathworks, Inc., 2009). Let $x_1 \ldots x_n$ be independent and identically distributed (iid) random variables with the cdf, $F(x)$. The empirical distribution function $F_n(x)$ based on sample $x_1 \ldots x_n$ is the step function defined by:

$$F_n(x) = \frac{\text{number of elements in the sample} \leq x}{n}$$

$$= \frac{1}{n}\sum_{i=1}^{n} I(x_i \leq x)$$

(5.10)

where $I(x_i \leq x)$ is the indicator of the event in parenthesis. Once the two bounding ECDFs are calculated, the 10^{th} and 90^{th} percentile points are obtained from each. These are shown in Figure 5.13 by asterisks.
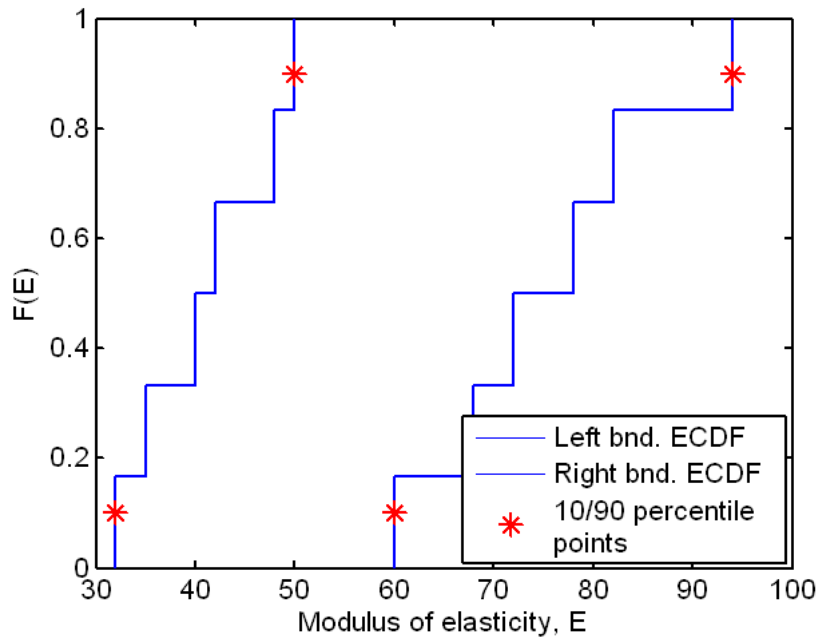
Figure 5.13. ECDFs for lower and upper bounds on *E* and 10th and 90th percentile points

A procedure to generate realizations of the parameters of a bounded Johnson distribution that is based on the experts'-given bounds shown in Table 5.2 and enforces the ECDFs shown in Figure 5.13 is described below:

1. Start by constructing the bounding ECDFs from the experts' bounds and identifying their corresponding 10^{th} and 90^{th} percentile points. This is shown in Figure 5.13. The percentile points are tabulated below.

Table 5.4. Lower and upper bounds on the 10th and 90th percentile points based on ECDFs of experts' given intervals

	Lower Bound	Upper Bound
10th Percentile point	32	60
90th Percentile point	50	94

2. Assume that the 10th and 90th percentile points follow a uniform distribution with limits given by the lower and upper bounds on the 10th and 90th percentile points (shown in Table 5.4). These distributions are shown in Figure 5.14.
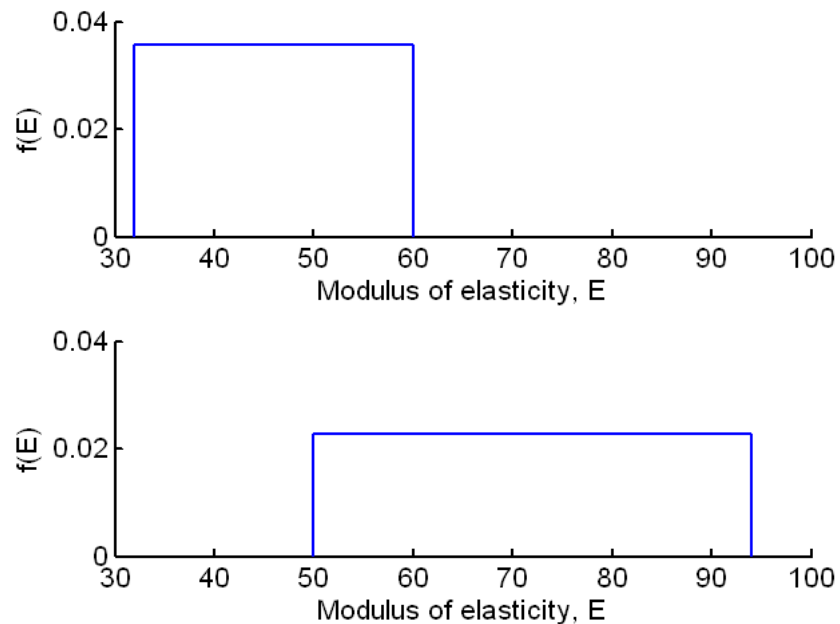
Figure 5.14. Uniformly distributed 10th percentile (upper graph) and 90th percentile (lower graph) with limits from ECDFs from experts' given bounds

3. From the distributions shown in Figure 5.14, sample n realizations of the 10^{th} and 90^{th} percentile values and denote these as:

$$\alpha = \begin{bmatrix} (\text{10th percentile})_i & (\text{90th percentile})_i \\ \vdots & \vdots \\ (\text{10th percentile})_n & (\text{90th percentile})_n \end{bmatrix} \quad i = 1...n \qquad (5.11)$$

4. Now, from the form of the bounded Johnson distribution given in Equation 5.8 and using the given interval data, the parameters λ and ξ can be established. These are the same as in the moment matching method:

$$\xi = 32$$
$$\lambda = 62$$

5. Now use the method of percentile matching and 1^{st} moment bounding to estimate remaining parameters, δ and γ. This is done using an optimization formulation as shown below:

find δ, γ

$$X = \xi + \lambda \cdot f^{-1}\left(\frac{Z-\gamma}{\delta}\right), \qquad \xi < X < \xi + \lambda$$

with $\qquad f^{-1}(z) = \dfrac{1}{\left(1+e^{-z}\right)}$ $\qquad\qquad$ (5.12)

subject to :
$$\alpha(i,1) = \hat{\alpha}_1$$
$$\alpha(i,2) = \hat{\alpha}_2$$
$$M_{1L} \leq \hat{m}_1 \leq M_{1U}$$
where
$$\hat{m}_1 = E(x - \mu_x)$$

where $\alpha(i,1:2)$ are the rows of α in Equation 5.12, $\hat{\alpha}_j$, $j = 1, 2$ and $\hat{m}_1$ are 10^{th}, 90^{th} percentile values and 1^{st} moment, respectively, found from sampled data taken from a bounded Johnson distribution with parameters λ, ξ, γ and δ and M_{1L} and M_{1U} are the lower and upper bounds on the 1^{st} moment respectively. The procedure was implemented in Matlab and a probabilistic description of the modulus of elasticity of foam is obtained and 30 realizations of the PDF of E are show in the figure below. The PDFs shown in Figure 5.15 show good coverage of the entire range of E and seem plausible realizations of the probability density for this variable.
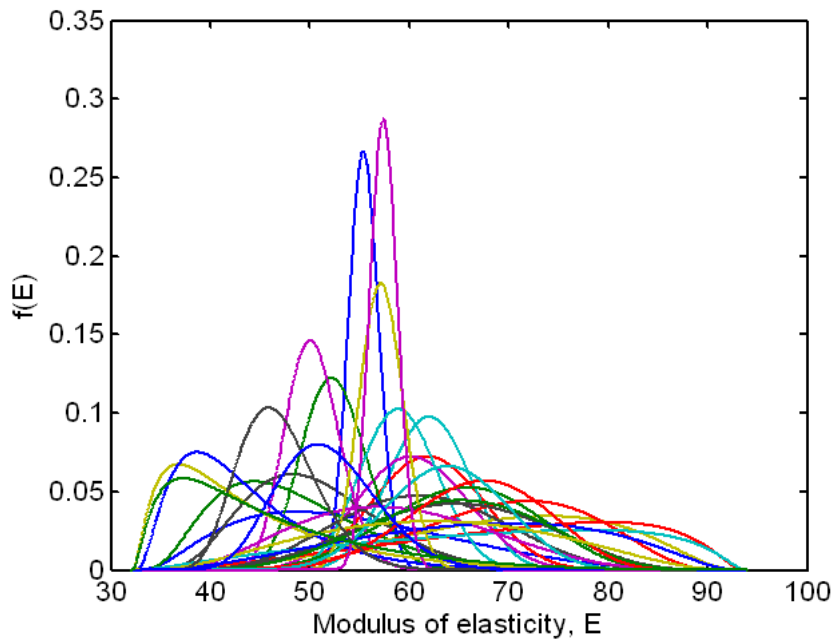
Figure 5.15. PDFs of foam modulus of elasticity using percentile matching method and level 0 data – These are the prior distributions of E

To verify that the algorithm is working correctly, a plot of the CDFs for each of the PDFs shown in Figure 5.15 are shown with their 10[th] and 90[th] percentile values along with the ECDFs and the corresponding bounding 10[th] and 90[th] percentile values. These quantities are shown in Figure 5.16.
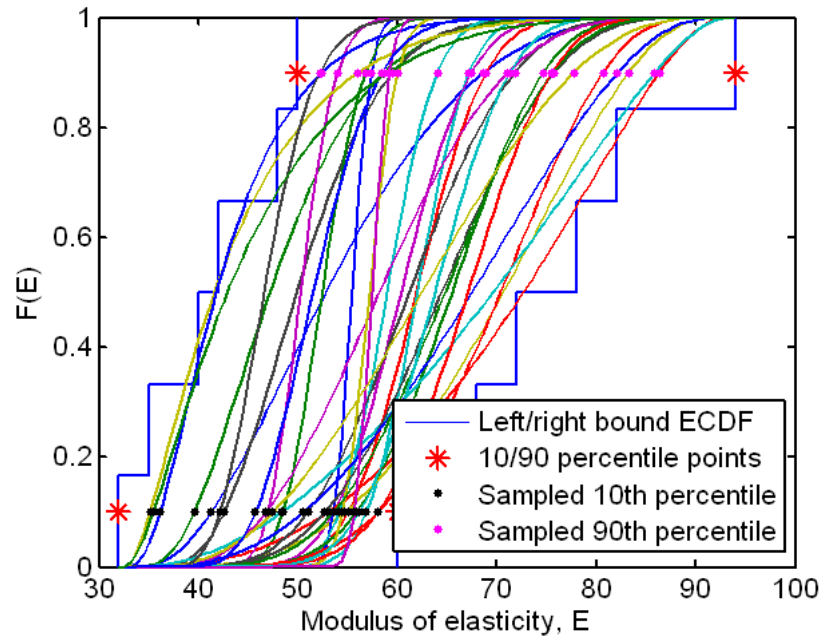
Figure 5.16. Estimated 10th and 90th percentile points from 30 generated CDFs

The final check on this methodology is to ensure that the first moments of the resulting PDFs fall within the calculated bounds on the first moments based on the expert's intervals and shown in Table 5.3. This is confirmed in Figure 5.17. Similarly to the method of moment matching, the 1[st] moments of the distributions obtained with the percentile matching method span the range defined by the bounds on the 1[st] moment obtained from the interval data. The realizations of PDFs for the variable E appear

plausible and fall within the prescribed criteria for this method (percentiles matched and 1[st] moments within the bounds).
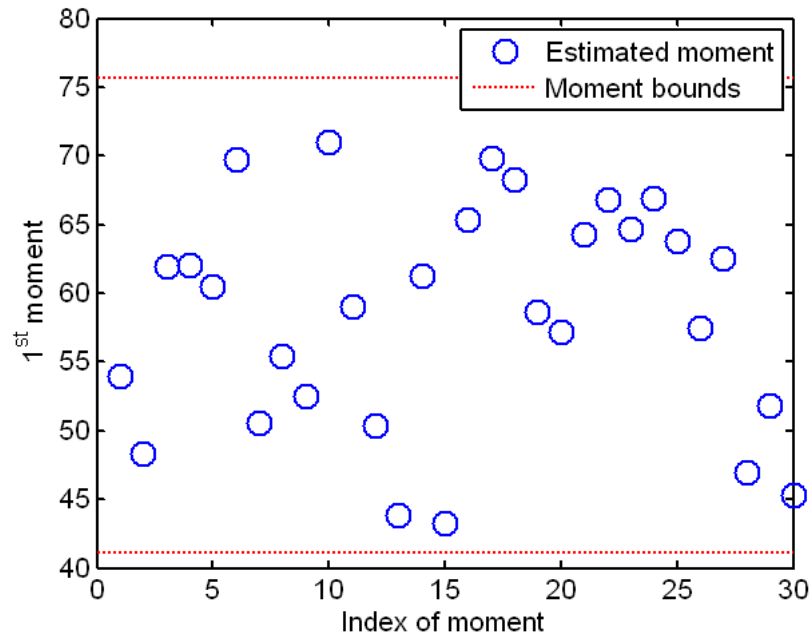
Figure 5.17. 1[st] moment of PDFs shown in Figure 5.15 and bounds on 1[st] moment from experts' given bounds

Similar to the previous method, the priors shown in Figure 5.15 are propagated through the Bayes network, updated via Markov chain Monte Carlo simulation, and used to obtain samples of the posterior distribution of E, and then the forward prediction of the system level response. This sequence is plotted in Figure 5.18 through Figure 5.19.
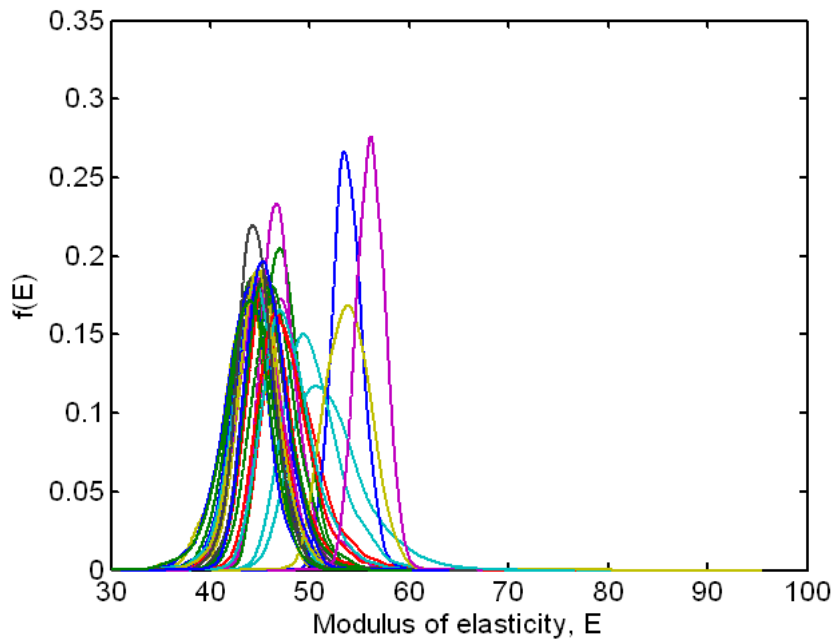
Figure 5.18. KDEs of posteriors of foam modulus of elasticity, *E*, based on level 1 and level 2 data (Note: Prior PDFs of *E*, based on level 0 data, are shown in Figure 5.15)
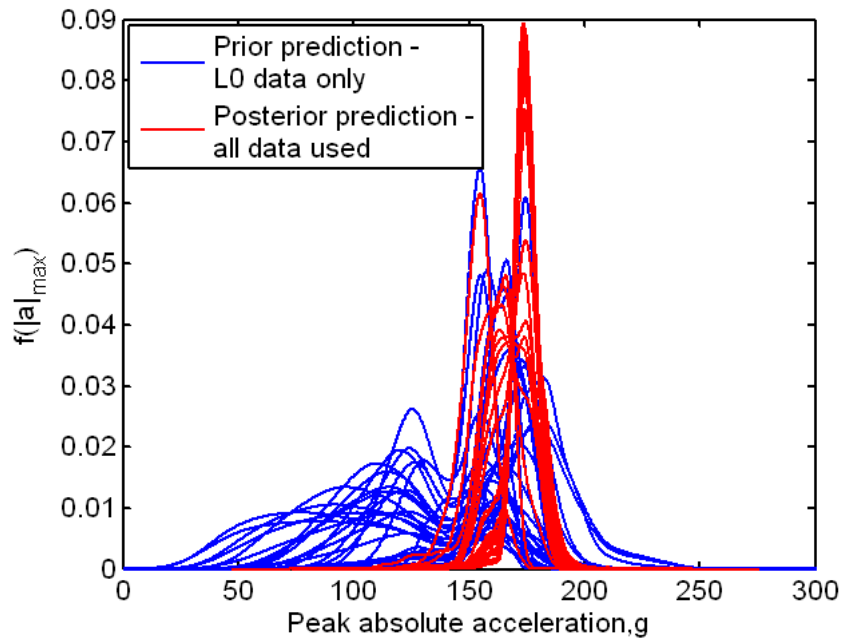
Figure 5.19. KDEs of peak acceleration from predicted responses of system level structures

114

When compared to the previous method (method of moment matching), it is clear that the priors are different from those generated before, and since the same data is used to update them, the resulting posterior distributions of modulus of elasticity show different variances and different means. This is a function of the location of the priors relative to the data used for updating and it is a reflection of the epistemic nature of the parameter (i.e. the true probabilistic form of the random variable is unknown). The system level prediction shown in Figure 5.19 demonstrate the effect of the varying posterior PDFs of the foam modulus of elasticity, E. The current method yields a larger scatter in the values of peak possible accelerations and the ranges of mean and variance of the resulting predictions are larger than those obtained using the previous method.

5.3.3   Parametric distribution functions from expert intervals

This method constructs parametric distributions functions using the bounded Johnson distribution as the framework. Furthermore, it models, individually, each of the expert-specified intervals. In other words, it treats the original problem of multiple intervals as six separate problems each defined by one of the intervals (as shown in Table 5.2). The six intervals are shown graphically in Figure 5.20.
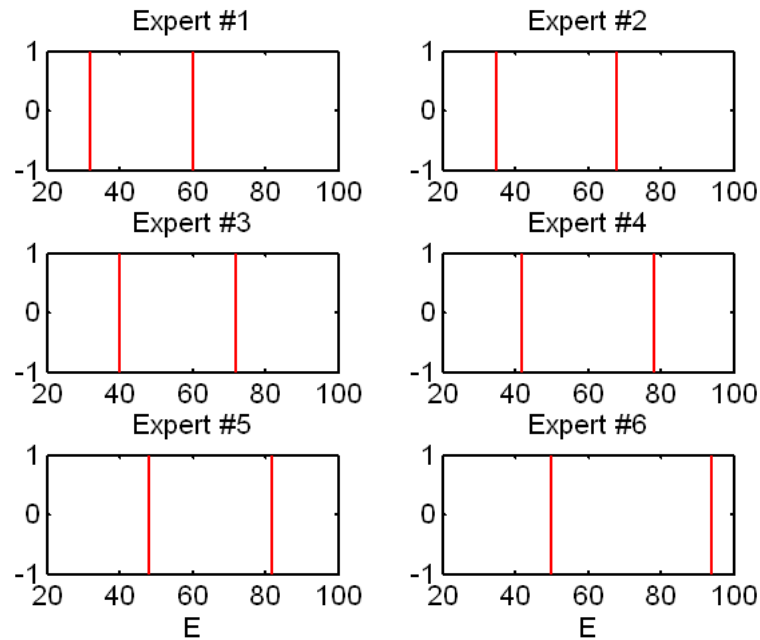
Figure 5.20. Expert-given intervals for values of modulus of elasticity of foam

Prior distributions for the foam modulus of elasticity are generated from each of the intervals shown in Figure 5.20 using a bounded Johnson distribution whose form is given by Equation 5.8. The parameters ξ and λ define the bounds within which this form of the Johnson distribution has realizations. The lower bound is ξ; the upper bound is $\xi + \lambda$. For example for the first expert, $\xi = 32$ and $\lambda = 28$. The parameters δ and γ are then chosen at random from within plausible ranges. (The ranges of delta and gamma are defined by examining the shape of the distribution function resulting from use of a set of parameters.) The ranges are: $\gamma = [-2 \ 4]$ and $\lambda = [0.33 \ 2.67]$. An analysis must define how many distributions are to be selected in each interval. One possibility is to select an equal number of distributions from each interval, thus assigning each expert an equal weighing. An example in which five distributions are generated from each interval is shown in

Figure 5.21. The 30 PDFs shown in the figure serve as priors of the foam modulus of elasticity.

As before, these prior distributions of modulus of elasticity are updated through the Bayes network to obtain estimates of the posterior distributions of modulus of elasticity; then each posterior PDF of modulus of elasticity is used in a forward prediction to obtain the PDF of peak acceleration of the system level structure. The PDF of priors and posteriors of modulus of elasticity, E, and peak acceleration response of the system level structure are shown in Figure 5.22 through Figure 5.24.
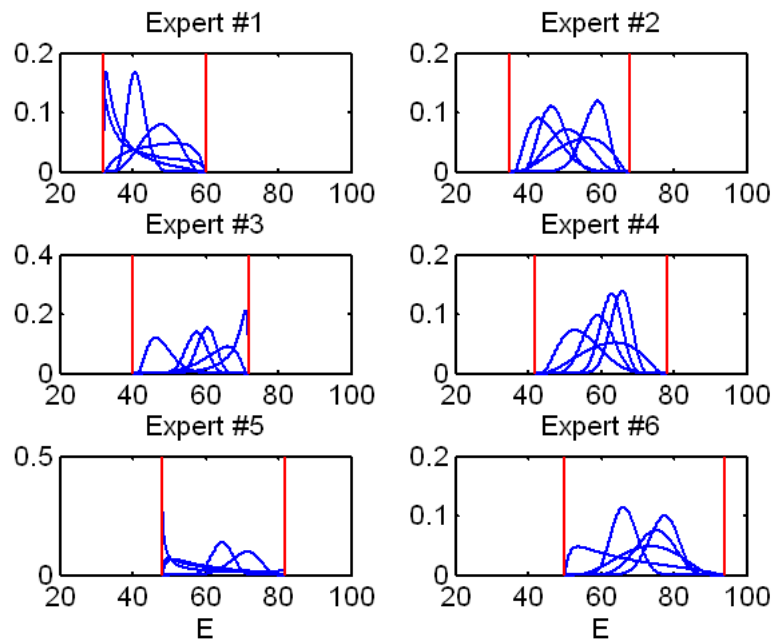
Figure 5.21. PDFs that serve as priors for Bayesian analyses. Five PDFs are defined for each expert-specified interval. This is considered equal weighing of expert-intervals
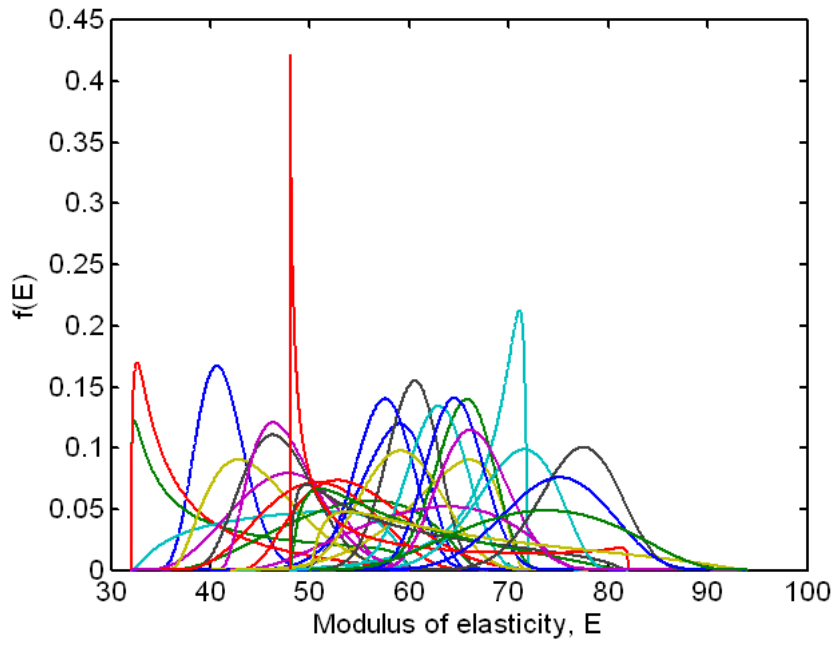
Figure 5.22. Collection of PDFs of priors of modulus of elasticity based on equal expert weighing

Figure 5.23. KDEs of posterior distributions of foam modulus of elasticity from expert equal weighing
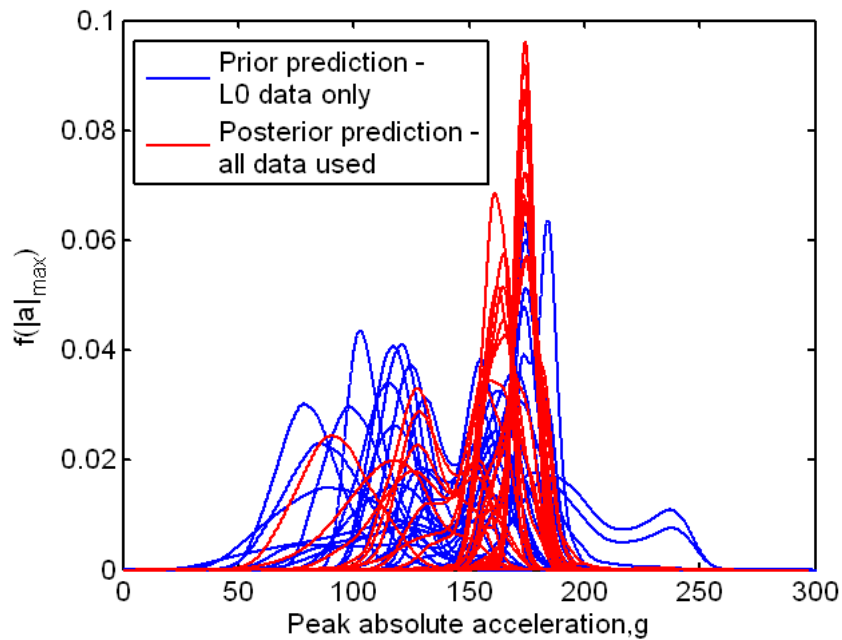
Figure 5.24. KDEs of peak acceleration responses from system level structural response predictions

One observation that can be made from Figure 5.24 is that the range of PDFs describing the system level response is much broader than the ranges observed using the previous methods. This is due, in part, to considering each of the experts as separate entities and not combining the information from all the experts. This method calls for each of the expert's opinions to be treated separately, thus, the resulting priors reflect each individual's range of values, only. The reason for the large range in posteriors is that some prior distributions start out far away from the available data at level 1 and 2 and thus the priors can only be updated to account for this data. Just to recall, the available data consisted of the following:

1. 45 joint experiments at level 1 and 27 joint experiments at level 2; energy dissipation was calculated for each experiment

2. 6 foam experiments at level 1 and 3 foam experiments at level 2; natural frequencies were calculated for each experiment.

Instead of equal weighing, the experts can be given different weights subject to their relative amount of experience, knowledge, credibility and/or past performance. For this research an arbitrary assignment of weights was defined. Here, weighing is interpreted in terms of a corresponding number of distributions in each of the expert's intervals. For example, the weighing (in terms of number of distributions in each interval) chosen for this work is shown in Table 5.5. The generated distributions are shown in their respective intervals in Figure 5.25.

Table 5.5. Relative weighing of experts (in terms of number of distributions in each interval)

Expert	Relative weight
1	1
2	12
3	5
4	3
5	7
6	2
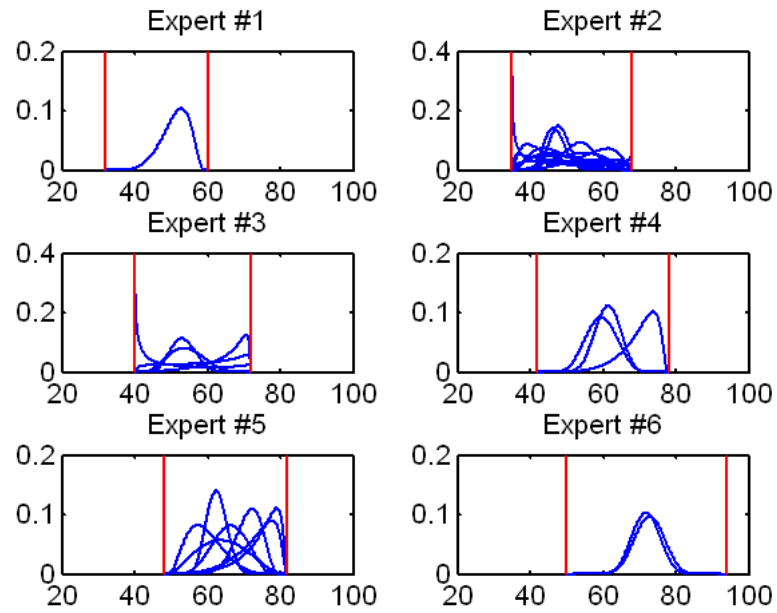
Figure 5.25. PDFs generated to represent expert-specified intervals. Weighting is non-equal.

The prior PDFs of foam modulus of elasticity are updated to obtain the KDEs of the posterior distributions of foam modulus of elasticity. Then each posterior KDE of foam modulus of elasticity is used to make a forward prediction of the PDF of peak acceleration response of the system level structure. The prior and posterior PDFs of foam modulus of elasticity and the PDFs of peak acceleration in the system level structure are shown in the sequence of Figure 5.26 through Figure 5.28. The KDEs of peak acceleration of system level response (shown in Figure 5.28) are very similar to those shown in Figure 5.24. The main difference is attributed to the fact that more prior distributions are generated for those experts whose opinions are weighted higher relative to the others. If some experts (experts 2 and 4), had priors that were closer to the data

used for updating, then the resulting posteriors and system level responses would tend to be more concentrated in the interval [160,190].
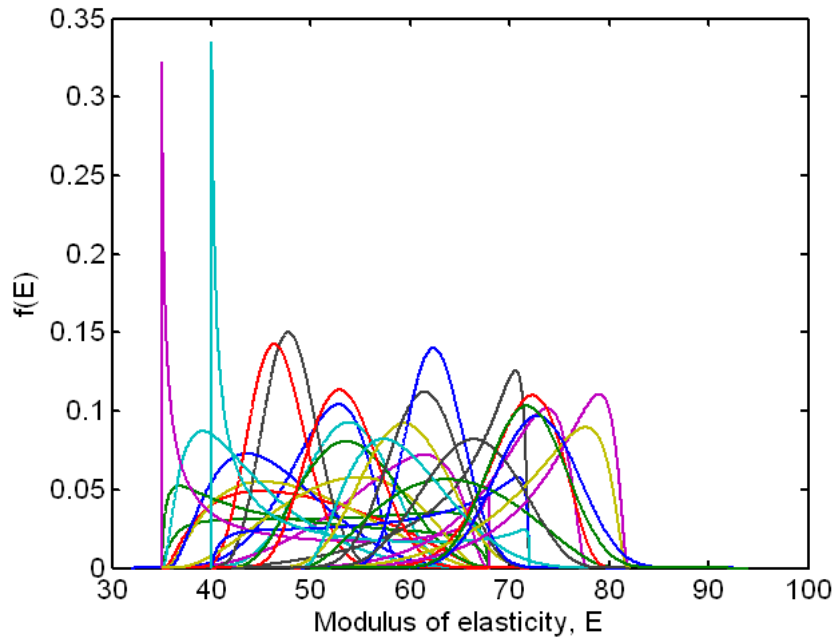
Figure 5.26. Collection of PDFs of priors of modulus of elasticity based on different expert weighing
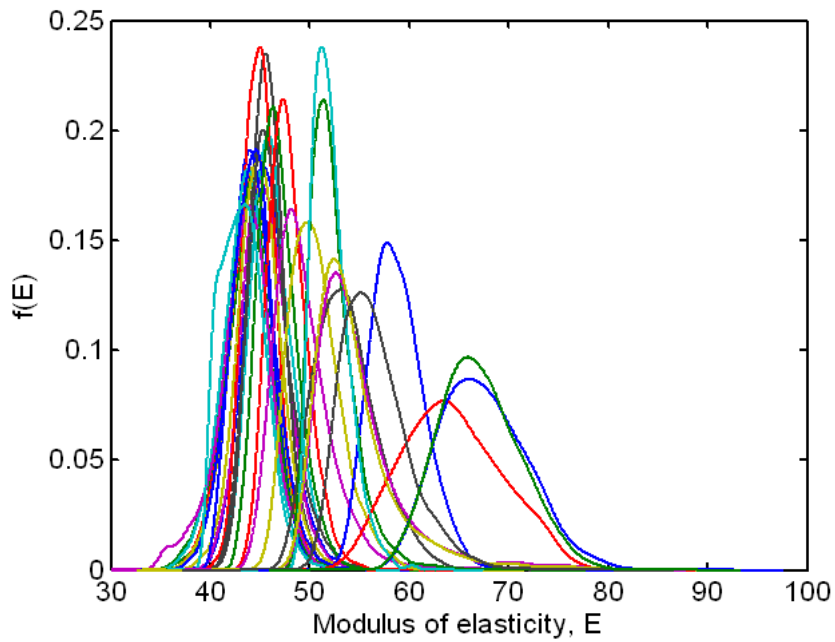
Figure 5.27. KDEs of posterior distributions of foam modulus of elasticity from expert different weighing
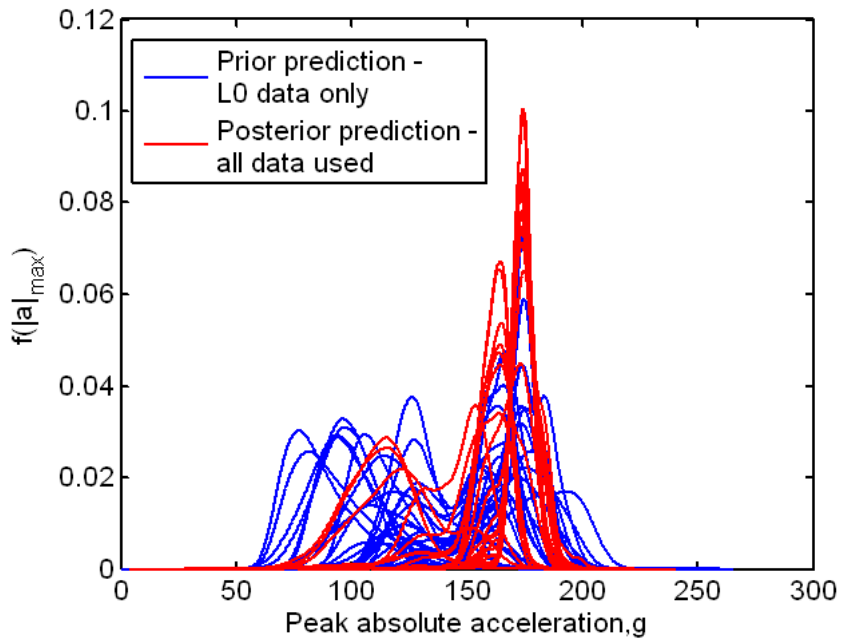
Figure 5.28. KDEs of peak acceleration responses from system level structural response predictions

5.3.4   Uniform distribution from each expert interval

This approach is conceptually similar to the one described above except that uniform distributions of foam modulus of elasticity are assumed for each expert interval. The uniform PDFs are shown in Figure 5.29.
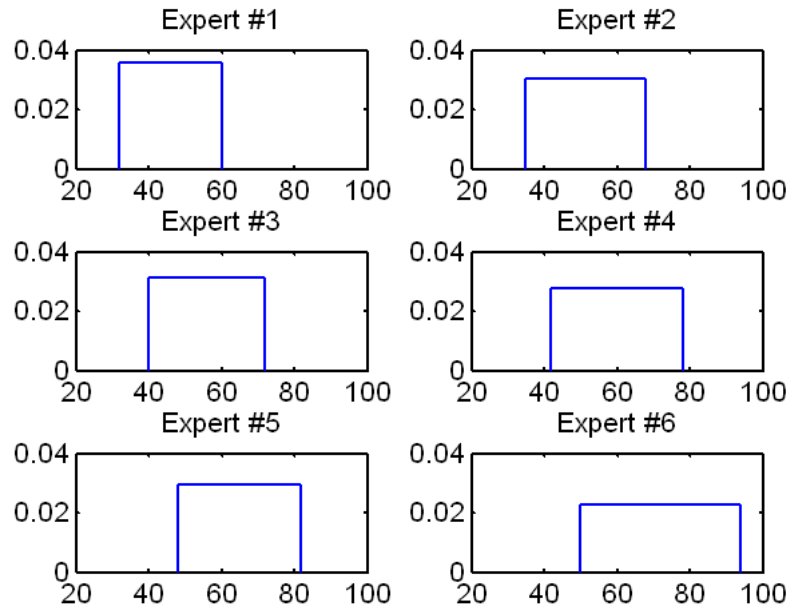
Figure 5.29. PDFs of expert intervals with a uniform distribution

For this particular case, there are only 6 prior distributions of modulus of elasticity to update into posterior PDFs of foam modulus of elasticity. Each of the posterior distributions was used to compute the KDE of peak acceleration response on the system level structure. A general observation is that the posterior PDFs of E and, in turn, the

forward predictions of PDFs of peak acceleration response are strongly influenced by the locations of the priors and the available data used to update the priors into the posteriors. It is not surprising that some of the PDFs of peak acceleration response group together around a central location of 180g; this is where the PDFs obtained using the other methods have also clustered.
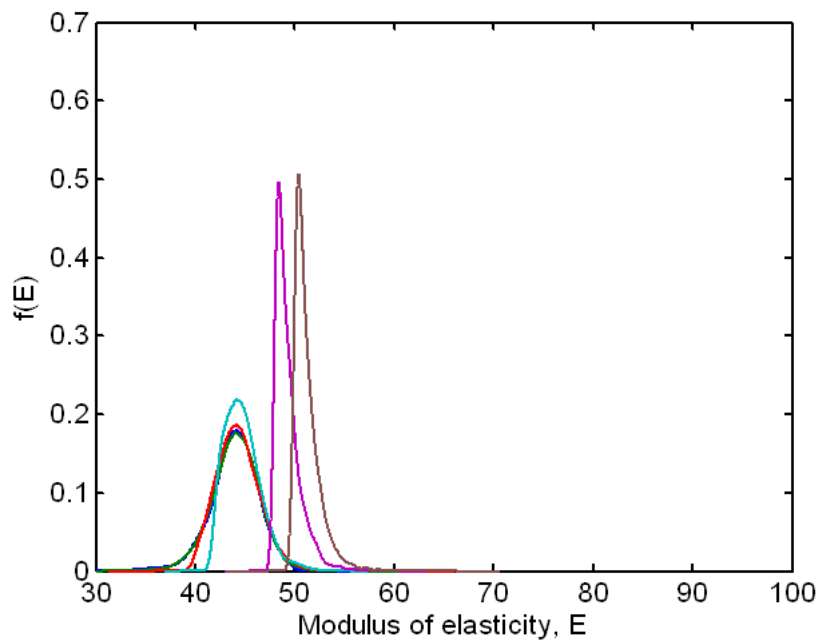
Figure 5.30. KDEs of posterior distributions of foam modulus of elasticity from 6 uniform priors
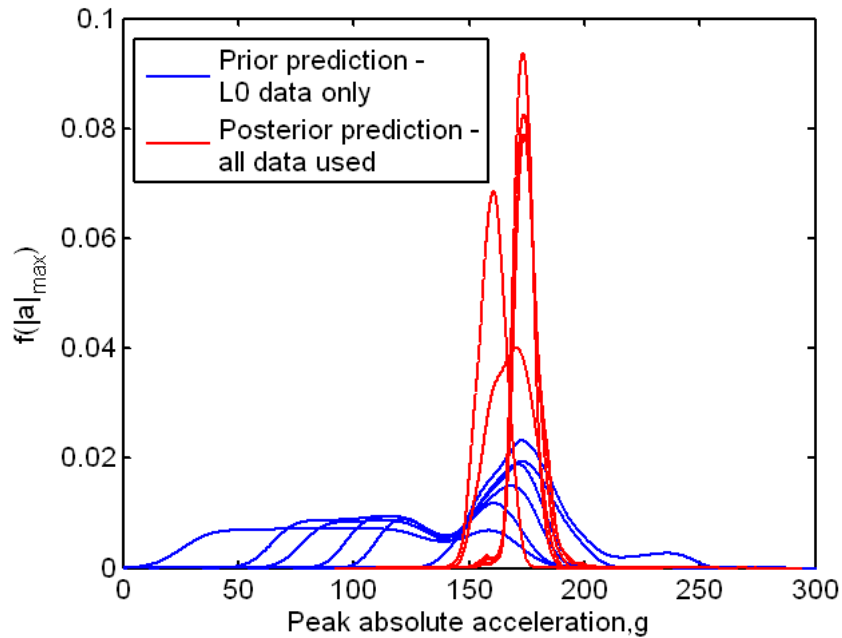
Figure 5.31. KDEs of peak acceleration responses from system level structural response predictions

5.3.5 One uniform distribution covering all the experts' range

The final method to incorporate the expert-specified information is to take the absolute minimum and maximum values given by experts and use this to define a uniform distribution. This defines a prior that encompasses all the ranges of values offered by the experts, but does not take into account the experts individual information. The PDF of this prior distribution is shown in Figure 5.32 and it provides a single posterior PDF of foam modulus of elasticity, which can be used to obtain one distribution of the peak acceleration of system level response. A family of distributions is obtained with the all other methods. This is analogous to performing an analysis where the modulus of elasticity is treated as an aleatoric variable and not an epistemic one. The

reason for performing this analysis is to determine the effect on the analysis with and without epistemic uncertainty included. The approach of selecting one prior distribution as in this case, is the most uninformed case since it uses the least amount of information to capture epistemic uncertainty. It uses a uniform distribution to encompass the entire range of values that are specified by the experts by setting the limits of the uniform distribution to the minimum and maximum of the experts' given intervals.
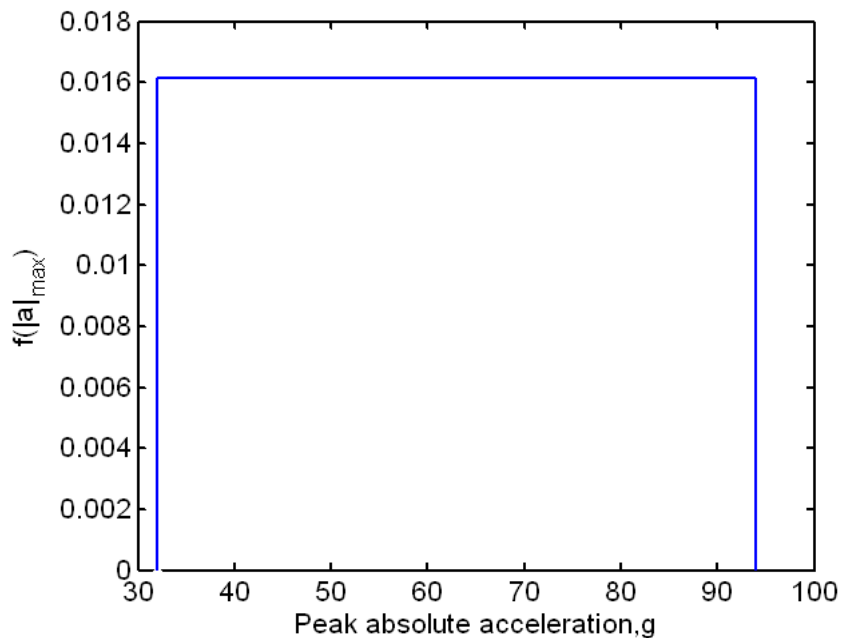
Figure 5.32. Uniform distribution between absolute bounds on expert-specified data – This is the prior distribution of *E*

Figure 5.33 and Figure 5.34 show the KDE of the posterior distribution of *E* and the KDE of the system level response. It is interesting to note that the resulting system level

response shown in Figure 5.34 is similar to the resulting KDE obtained in Chapter 4 when all random variables were treated as sources of aleatoric uncertainty.
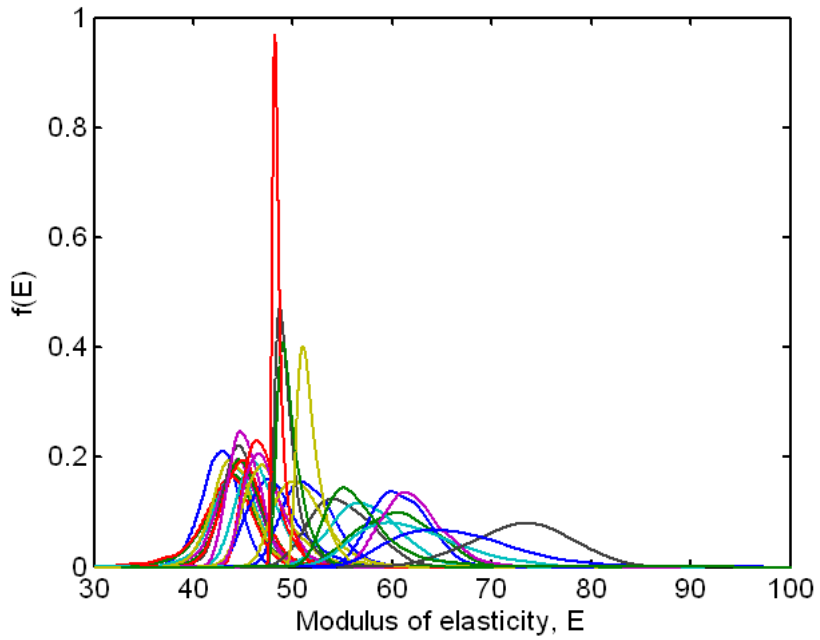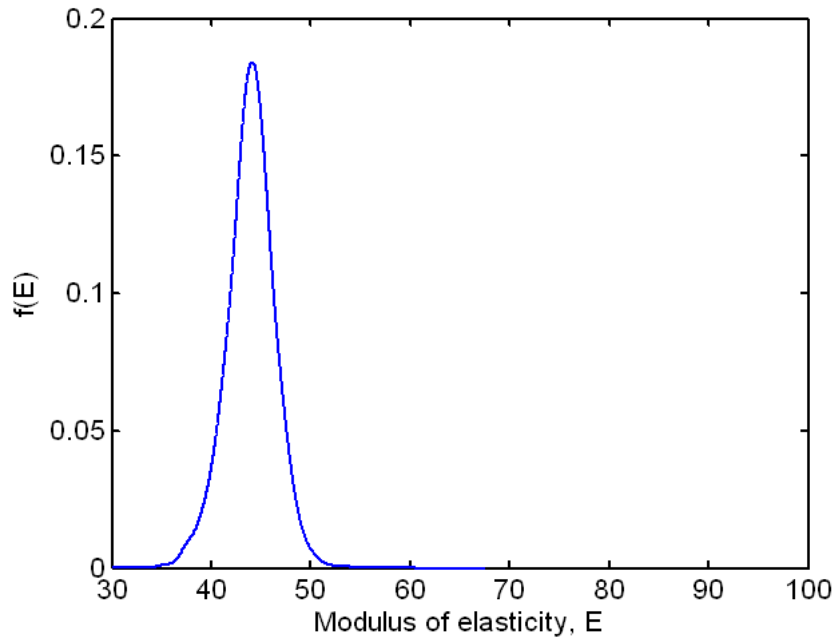

Figure 5.33. KDEs of posterior distribution of foam modulus of elasticity
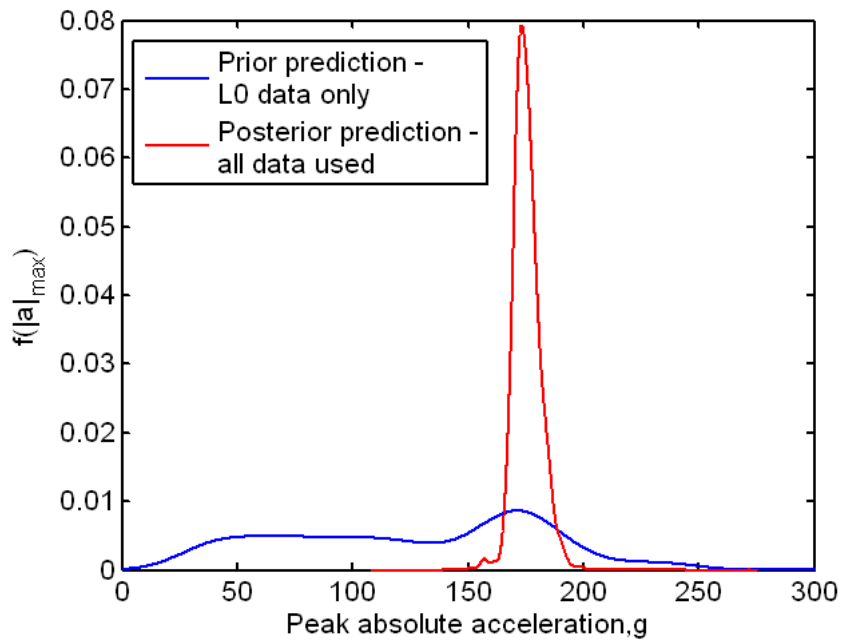
Figure 5.34. KDE of peak acceleration responses from system level structural response prediction

5.4 Comparison of Results

In the following two sections, the results presented in Section 5.3 are compiled and compared in order to make some general statements regarding the suitability of each of the methods to incorporate epistemic uncertainty into a Bayesian analysis of a hierarchical model. Section 5.4.1 presents results which arise due to UQ method uncertainty. This section summarizes the results for all the methods shown in Section 5.3. Section 5.4.2 shows the model error comparison for four of the methods described in Section 5.3. The errors are also compared to those presented in Chapter 4 where only aleatoric uncertainty was considered.

5.4.1 UQ Method Uncertainty - Results

This section presents the effect of the uncertainty quantification methodology used to model the modulus of elasticity of foam on the system level prediction. First, the prior distributions for the various analysis methods are plotted in Figure 5.35 (Note that the uniform priors are not shown.) To facilitate visualization of all the results, the cumulative density functions (CDF) are plotted instead of the PDFs. Although in this figure it is difficult to examine small features of individual distributions, it is most helpful to look at some global characteristics of the CDFs. One of the most salient features is the shape of the CDFs from the moment matching technique. These CDFs are consistent with distributions that are not unimodal, are highly skewed and are relatively flat. These PDFs can be observed in Figure 5.9.

The next interesting feature of this plot is the range of values that these priors cover. Of course, the extreme values are those corresponding to the limits of the expert opinion (i.e. from Table 5.1 the values are 32 and 94). It is noted that some distributions start and end well away from these endpoints, and they are all reasonable quantifications of epistemic uncertainty.
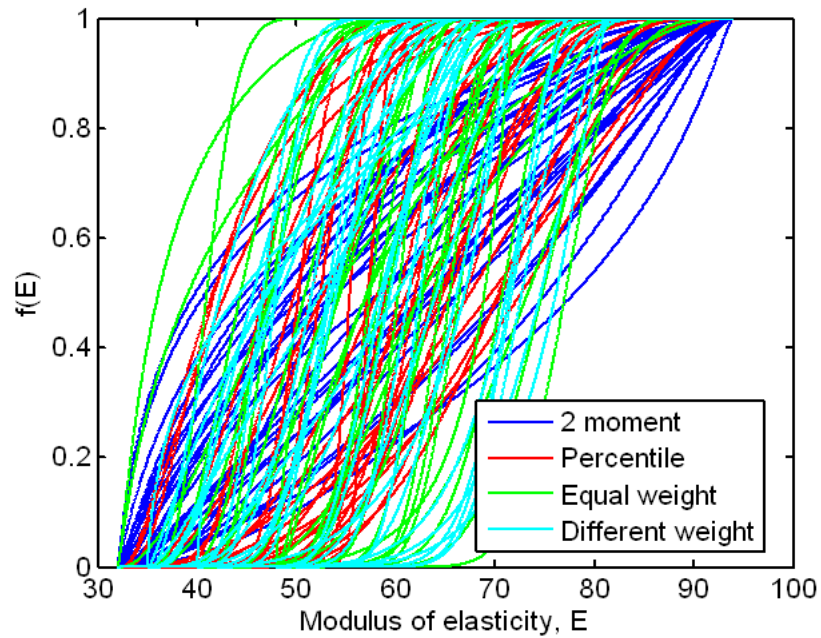
Figure 5.35. Compilation of all prior CDFs of foam modulus of elasticity from different methods

Figure 5.33 shows the CDFs of the posterior distributions of the foam modulus of elasticity for all methods. The curves show, clearly, the effect of the data used for updating relative to each of the priors. In general, when the prior distributions have support over the range where data is present, the prior moves toward the data during the updating procedure, and the posterior reflects the characteristics of the prior and the likelihood function. When the prior distribution is far away from the data, the PDF of the posterior distribution may resemble a delta function, thus indicating little or no support for the prior distribution in the likelihood function. In general, the CDFs shown in Figure 5.36 show that all prior distributions have some degree of data support.
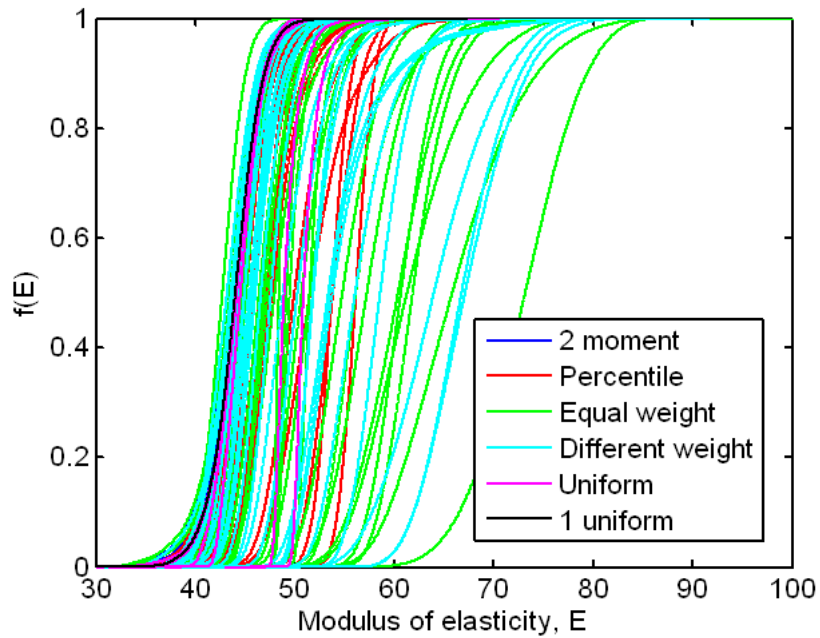
Figure 5.36. Posterior distributions for all methods

The next comparison considers the distributions of system level peak acceleration response predictions. These are shown in Figure 5.37. These CDFs reflect not only the effect of the epistemic variable (foam modulus of elasticity) but also the effects of the aleatoric variables (the joint model parameters, k_{lin}, k_{non} and npow). The system level response is an extrapolated quantity and no experimental data has been used for updating (although it could be easily incorporated if it were available).
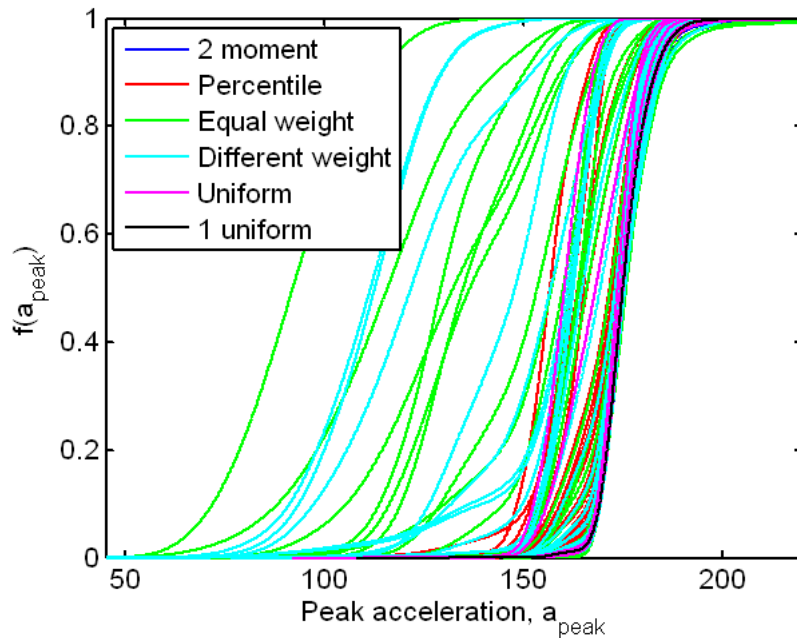
Figure 5.37. System level predicted responses for all methods

With the exception of some of the equally weighted and differently weighted expert CDFs, the majority of the system level responses cluster around values between 140g to 200g.

The next set of plots compares the results obtained in this chapter with those obtained in Chapter 4. This is a comparison of the effect of both epistemic and aleatoric uncertainty present in an analysis relative to when only aleatoric uncertainty is considered. Figure 5.38 shows the prior distributions of E. As it is apparent from the figure, the distribution of this variable has a much greater spread when interval data is considered but it is also constrained to a certain range. This is to be expected as the intervals given by the experts need to be preserved.
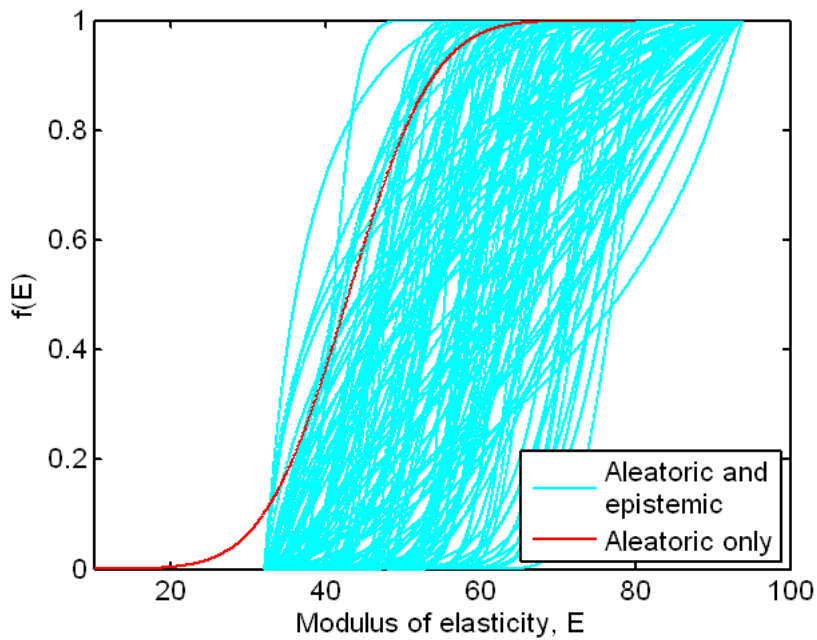
Figure 5.38. Comparison of prior distributions of *E* - Aleatoric/Epistemic and Aleatoric only

Figure 5.39 shows the posterior distribution of *E* once all the available data is used to update the Bayes network. It is interesting to note that the CDF of *E*, when only aleatoric uncertainty is considered, falls towards the lower end of the range defined by the CDFs obtained when both epistemic and aleatoric uncertainty is included. This says that the possible range of *E* is much greater when epistemic uncertainty is included in the analysis and also due to the different method used to quantify uncertainty due to interval data.
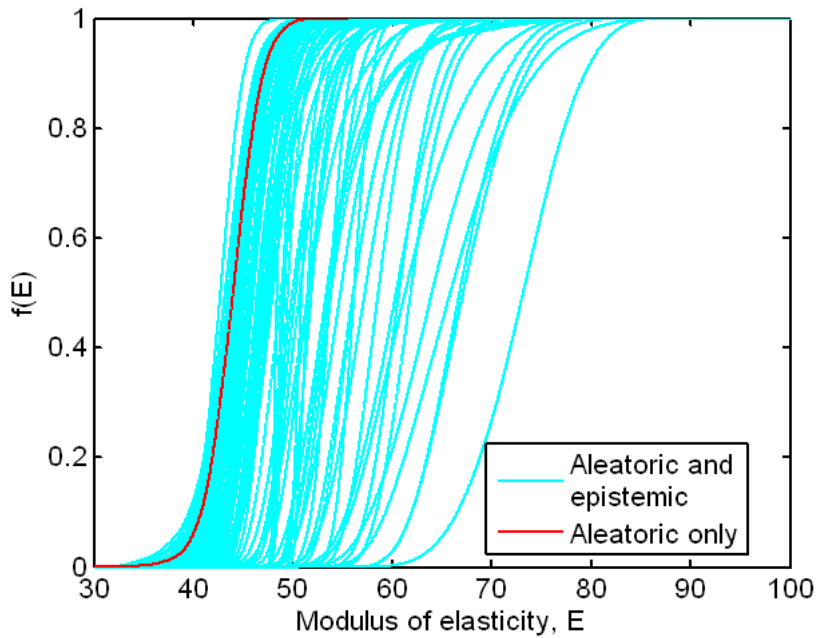
Figure 5.39. Comparison of posterior distributions of E - Aleatoric/Epistemic and Aleatoric only

Finally, Figure 5.40 shows the comparison of the posterior distribution of the system level response with and without epistemic being considered. Consistent with the results shown in Figure 5.39, the system level CDF when only aleatoric uncertainty is included falls on the upper end of the range of CDFs (as noted in Chapter 3, there is an inverse relationship between E and the system level response). The presence of epistemic uncertainty has an effect to expand the range of possible system level responses and thus it accounts for possibilities that are not included if the parameter E is treated as an aleatoric variable.
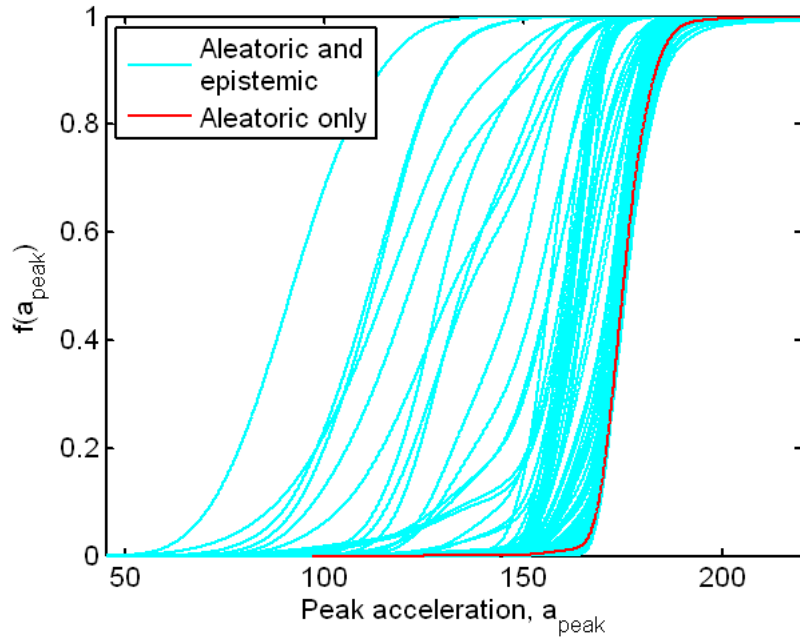
Figure 5.40. Comparison of posterior distributions of system level response - Aleatoric/Epistemic and Aleatoric only

5.4.2   Model Errors - Results

This section presents the statistics of model error terms, denoted as $\varepsilon_1^j, \varepsilon_2^j, \varepsilon_1^f$ and ε_2^f (see Figure 4.1) and compares four of the methods examined in Section 5.3: moment matching, percentile matching, equal and unequal weighing of experts. The results obtained when interval data is present are also compared with the results obtained when treating parametric uncertainty as aleatoric uncertainty (as was presented in Chapter 4). Note that the prior distribution for the error terms is the same as what is shown in Equations 4.6 and 4.7.

The first set of plots, shown in Figure 5.41, presents the posterior distribution of the error terms associated with the moment matching method of treating interval data.
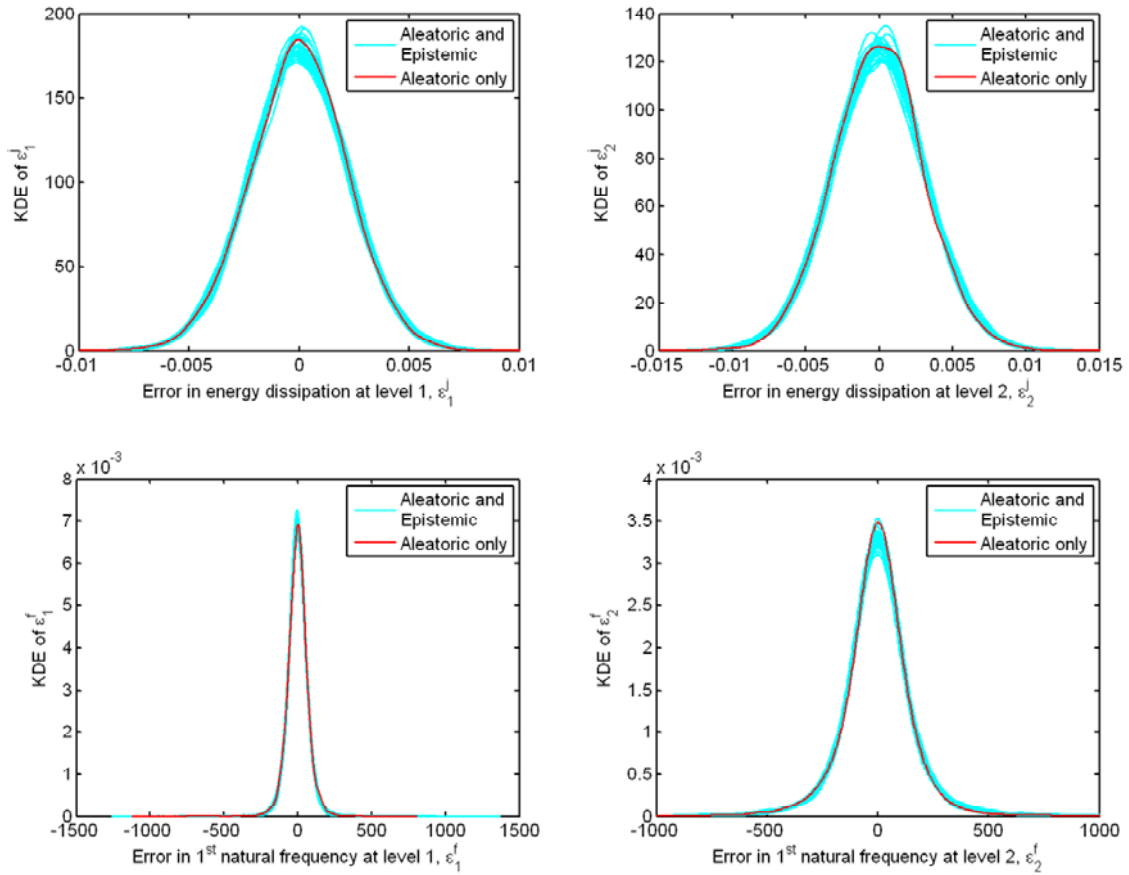
Figure 5.41. KDE of error terms - using moment matching method

One interesting feature is that the results when both epistemic and aleatoric uncertainty is considered appear to include those obtained when only aleatoric uncertainty is present; this is expected. The collection of KDEs for the epistemic/aleatoric case are due to the multiple realizations of the prior distributions of E.

Figure 5.42 shows the error terms when the percentile matching method is used to obtain the parameters of the Johnson distribution used to model interval data.
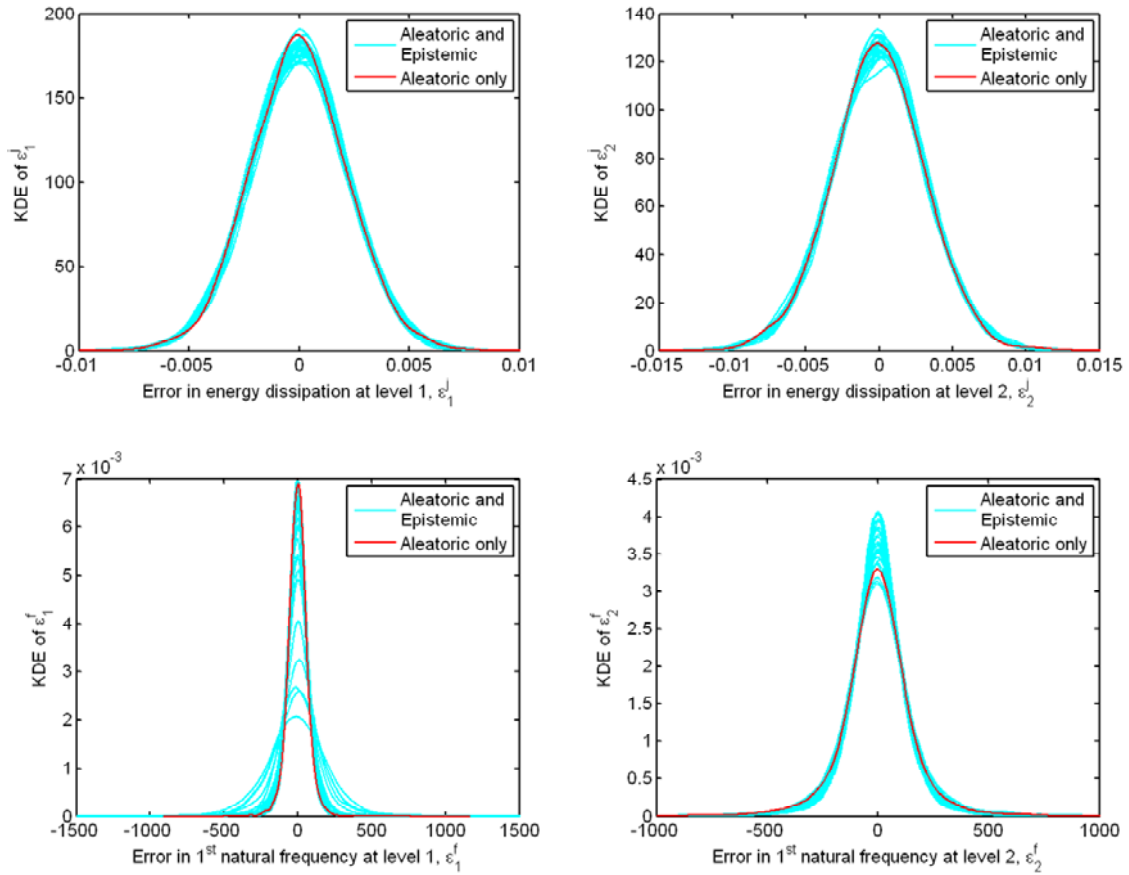
Figure 5.42. KDE of error terms - using percentile matching method

An interesting effect is shown in the error of the 1st natural frequency of the foam level 1 hardware. The variance changes for each realization of the prior distribution of the modulus of elasticity and in general it is higher in all cases relative to the aleatoric only case. This says that the error in this metric is highly sensitive to the variation in the prior of E. The other error terms do not show this effect.

Figure 5.43 and Figure 5.44 show the error terms for the non-aggregating methods used to model interval data: equal and unequal expert weighing. Similar to the results

shown for the percentile matching methods, the variance change in the error of the 1st

natural frequency of the foam level 1 response measure is noticeably different for both

cases and in general higher when compared to the aleatoric only case.
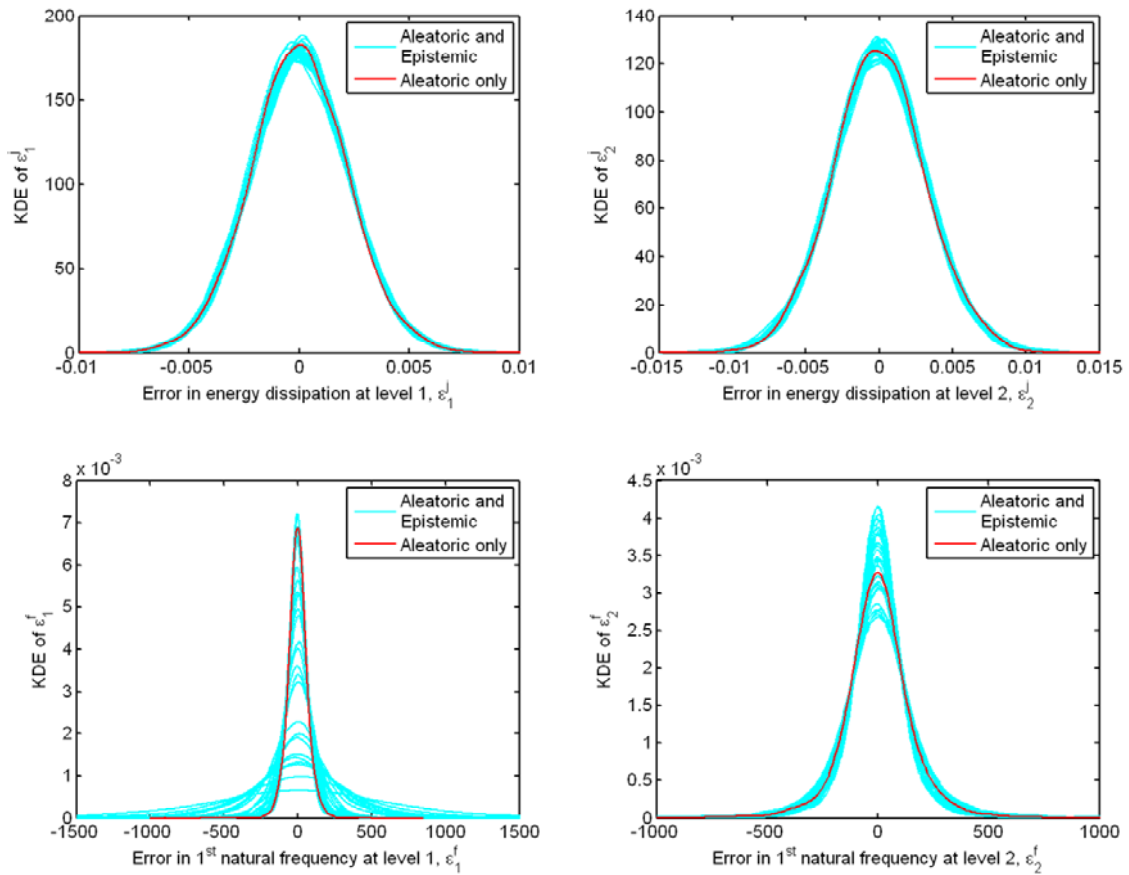
Figure 5.43. KDE of error terms - using equal weighing of experts method
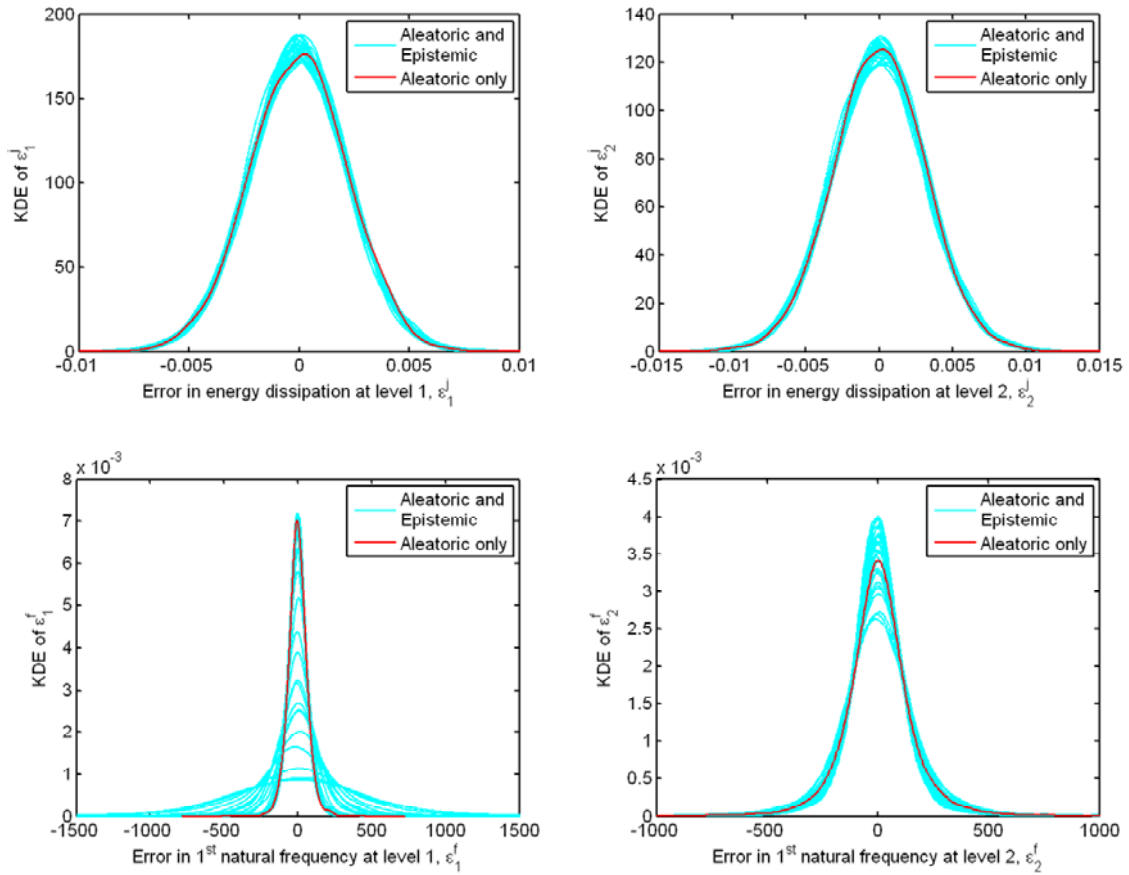
Figure 5.44. KDE of error terms - using unequal weighing of experts method

5.5 Conclusion

In this chapter, the treatment of interval data using a probabilistic approach was presented. Results obtained from propagating both epistemic and aleatoric uncertainty using a Bayes network were shown. In the literature, interval data has been addressed with non-probabilistic methods but in this research a decision was made to pursue a probabilistic approach to enable the implementation of Bayes networks and eventually, a reliability based approach to address the main focus of this research which is risk informed decision analysis. Different approaches to treat interval data when used to

describe a model parameter lead to multiple PDFs of the system level response. This is an example of model form uncertainty. This arises from the treatment of the so-called parametric uncertainty when the modulus of elasticity is given in terms of an interval. Both of these are sources of epistemic uncertainty

An interesting factor reflected in the analyses in this chapter is the effect of UQ method uncertainty in incorporating interval data. This comes to light due to the choice of methodology used to describe the modulus of elasticity given bounds on the parameter. In this chapter, six different ways to model interval data are presented. Using the methodologies described in this chapter, the resulting uncertainty at the system level was quantified and comparisons among the various techniques were made. It was observed that some techniques yield results that were very consistent among each other while others had a much larger range of values of the system level response. This was the result of the location and shape of the prior distributions and the available data used for updating. At this point, there is no attempt to rank these methods or suggest which ones perform better. This would be later revisited in the context of a decision making framework.

Also shown in this chapter is a comparison of the results obtained in Chapter 4 where only aleatoric uncertainty was included in the analysis. Both the results from the modulus of elasticity and the system level response were compared. It was observed that when epistemic uncertainty is present the range of possible values for these quantities was increased and in general the results from the aleatoric only analysis tends to fall on the extreme of the distributions obtained when both aleatoric and epistemic uncertainty are considered.

A comparison was also made using the model error terms present in the Bayes network. When both epistemic and aleatoric uncertainty is present, there is a collection of model errors that arises due to treating the interval data probabilistically. Two comparisons were made. One was the effect of including epistemic uncertainty which shows a minimal effect on some of the error distributions relative to the aleatoric only analysis. The other effect was due to UQ method uncertainty. In this case, the effect was clearly different based on the choice of method used. Except for the results of using the moment matching method, the other methods showed a significant change in the variance of the errors associated with the level 1 and 2 foam. This is an interesting result and it speaks to the sensitivity of this type of uncertainty in the overall analysis.

CHAPTER VI

RELIABILITY-BASED QMU

6.1 Introduction

A central step in the process of risk informed decision analysis (RIDA) is the Quantification of Margins and Uncertainty (QMU). QMU methodology has been, for the most part, proposed at the conceptual level and not much literature is available describing actual applications. In this regard, this dissertation looks at a few candidate implementations of this methodology and selects the most appropriate one to an actual problem developed at Sandia.

Based on the current literature, one possible implementation is described in Diegert et al. (2007). The QMU measure presented in this work is the confidence factor, (*CF*), and it is defined as:

$$CF = M / U \tag{6.1}$$

where M is the margin between the system behavior and the required performance measure, and U is defined as the *uncertainty* on the operating region and is assessed through modeling, testing, expert judgment or some combination of the three. In practice, U captures both aleatory and epistemic uncertainties. To ensure consistency across applications, it is recommended that the margin M be defined in terms of the difference of median values for assessed and threshold distributions and that the uncertainty U be defined in a manner to convey "high confidence" in the context of a specific application. If the assessment of U is rigorous, then it is sufficient that $CF > 1$ to ensure that the

reliability is "ONE" with high confidence. In practice, however, it is prudent to demand some robustness to unknown unknowns or to assessments lacking rigor in the modeling processes. Consequently, it is likely that some issues will require additional attention if *CF* is too close to unity (Diegert et al., 2007). An application paper based on the work by Diegert et al. is described in Pepin et al. (2008). This paper describes a step by step procedure to perform a QMU analysis based on Diegert et al.'s approach. The essence of this paper is that a unique way to implement QMU is not currently available and therefore opportunities abound to suggest techniques that satisfy the requirements for a QMU analysis.

A second approach to implement QMU is based on the risk-based decision methodology of Jiang and Mahadevan (2007) which is quite suitable for the analysis done in this research. This work was proposed to answer a model validation assessment question but it could be extended to address a QMU type question. The result of the risk-based methodology presented by Jiang and Mahadevan (2007) is the use of Bayesian hypothesis testing and the Bayes factor as the comparison metric. The proposed techniques should be valid approaches to implementing QMU as long as they satisfy the Kaplan and Garrick risk triplet plus the credibility component as described in Chapter 2. As a matter of fact, the fourth component, credibility, is the only one that relates to a quantitative assessment of the system performance and it is the one that can differ in the implementation. The other three components are more or less defined by the system being evaluated (i.e. the weapons system), the use environment and the ultimate objective of the system. For the purpose of this research, it will be assumed that these remain

constant for a particular system so this research will only propose a method to address the credibility issue.

Another possible approach to the treatment of QMU, in particular, one that treats interval data with probability theory, is to consider techniques used in the field of reliability using Bayes networks. The basic concept of reliability analysis can be found in Haldar and Mahadevan (2000). Estimation of systems reliability using Bayes networks dates back to 1988, when it was first defined in Barlow (1988). The idea of using Bayes networks in systems reliability analysis has gained acceptance because of the simplicity it allows in the representation of systems and the efficiency for obtaining component associations. Recently, Bayes networks have found applications in fault detection systems (Jensen, 2001) and general reliability modeling (Bobbio et al., 2001). Bayes networks have been developed for reliability estimation for specific systems. Gran and Helminen (2001) provide a Bayes network for nuclear power plants and introduce a hybrid method for estimating the reliability of the plant. Wilson and Huzurbazar (2006) showed using a simple two level system that it is possible to relate multiple levels of complexity to a system reliability analysis within a Bayesian context.

The reliability analysis in this research will be coupled to the use of Bayes networks to probabilistically combine the information available at multiple levels of complexity leading up to the system level. The major reason for using this methodology is the need to ultimately cast the decision making problem in a probabilistic framework and provided decision makers with a probability of occurrence rather than just a deterministic value. Still, the major complication and the real crux of the problem is the way that interval data is handled, and a key question concerns how much uncertainty is

added due to the choice of treating the interval data in a probabilistic way. This is something that can be answered only if a comparison with another approach to interval data treatment is performed. In the following section, a method to calculate the probability of failure based on the system level results is presented and a comparison using the various methods used to model interval data is also shown.

6.2 Methodology

As stated in the introductory section, the approach taken to implement QMU in this research is one that is based on reliability-type analysis. In this context the calculation of the probability of failure (or the probability of exceeding a given threshold) needs to be performed using the results obtained from the Bayesian network described in Chapter 4, which results from the inclusion of epistemic and aleatoric uncertainty as shown in Chapter 5. The methodology applied here makes extensive use of kernel density estimators (KDE) whose form was presented in Equation 4.4 of Chapter 4, to facilitate the computation of the probability of failure. The derivation of a formula to obtain this is shown below. (Key references to basic probability theory and reliability concepts can be found in Ang and Tang (1975), Haldar and Mahadevan (2000)).

To begin, let S be a random variable describing the response quantity of interest obtained from a model of the system. In the example problem, S describes the absolute peak acceleration response of the system, given a particular input. Now, let R be a random variable denoting the design threshold of the system given the same input to the system as the one used for the model. The values of R are usually obtained from experiments, from historical data, experts or a combination of all these. The key is that

146

this is the threshold at which the system of interest is assumed to fail. Let Z denote the margin of safety against failure, then,

$$Z = R - S \tag{6.2}$$

The probability of failure is the given by:

$$p_f = P(Z \le 0) = P(R - S \le 0) = P(R \le S) \tag{6.3}$$

For this research, the random variable S is known only through a collection of samples, s_j, $j=1...n$, which are obtained from the Bayes network following updating and the probability distribution of R is assumed to follow a normal law with mean μ_R and variance, σ_R^2. The probabilistic characterization of S is an approximation to its PDF that uses the kernel density estimator (KDE) defined as (This is a repeat of Equation 4.5.):

$$\hat{f}_S(s) = \frac{1}{n} \sum_{j=1}^{n} \frac{1}{\sqrt{2\pi}\varepsilon} \exp\left[-\frac{1}{2\varepsilon^2} (s - s_j)^2 \right] \tag{6.4}$$

where the kernel used in the KDE is the Gaussian kernel with standard deviation, ε. (The value of ε can be optimized based on the sample standard deviation of the s_j, $j=1...n$, and the number of data, n). The random variable R is assumed Gaussian in this research (The assumptions is that a predetermined requirement has specified this.) Its PDF is given by:

$$f_R(r) = \frac{1}{\sqrt{2\pi}\sigma_R} \exp\left[-\frac{1}{2\sigma_R^2} (r - \mu_R)^2 \right] \tag{6.5}$$

Given the PDFs in Equations 6.4 and 6.5, we seek the CDF of Z,

$$F_Z(z) = P(Z \le z) = P(R - S \le z) \tag{6.6}$$

Assume the random variables R and S are independent random variables; this is a reasonable assumption since there is no reason to believe that the system level response and its threshold are related. Therefore,

$$F_Z(z) = \int\limits_{r-s\leq z} dr \int ds \, f_R(r) \, f_S(s) \tag{6.7}$$

The region over which the integral is evaluated is shown in Figure 6.1.
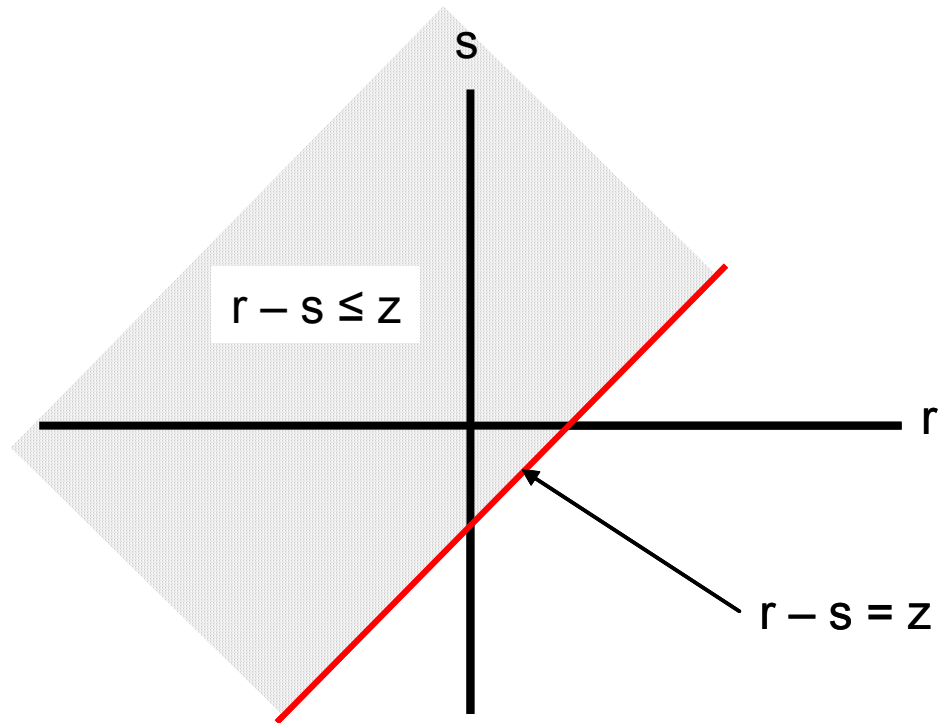
Figure 6.1. Failure region

The integral in Equation 6.7 can be rewritten (reflecting the region shown in Figure 6.1)

$$F_Z(z) = \int\limits_{-\infty}^{\infty} dr \int\limits_{r-z}^{\infty} ds \, f_R(r) \, f_S(s) = \int\limits_{-\infty}^{\infty} f_R(r) \, (1 - F_S(r-z)) dr \tag{6.8}$$

or

$$F_Z(z) = \int_{-\infty}^{s+z} dr \int_{-\infty}^{\infty} ds \, f_R(r) \, f_S(s) = \int_{-\infty}^{\infty} F_R(s+z) \, f_S(s) \, ds \qquad (6.9)$$

The PDF of Z is the derivative of the CDF,

$$f_Z(z) = \frac{d}{dz} F_Z(z) \qquad (6.10)$$

To evaluate the right hand side of Eq. 6.10, we need the derivative of an integral, which can be obtained by Leibnitz rule. The PDF of Z when derived using Equation 6.8 is (Paez, 2009):

$$f_Z(z) = \int_{-\infty}^{\infty} f_R(r) \, f_S(r-z) \, dr \qquad (6.11)$$

Substitute Equations 6.4 and 6.5 into Equation 6.11 to obtain,

$$f_Z(z) = \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}\sigma_R} \exp\left[-\frac{1}{2\sigma_R^2}(r-\mu_R)^2 \right] \frac{1}{n}\sum_{j=1}^{n} \frac{1}{\sqrt{2\pi}\varepsilon} \exp\left[-\frac{1}{2\varepsilon^2}(r-z-s_j)^2 \right] dr \qquad (6.12)$$

After simplifying Equation 6.12, the following expression is obtained:

$$f_Z(z) = \frac{1}{\sqrt{2\pi}\sigma_R n} \int_{-\infty}^{\infty} \exp\left[-\frac{1}{2\sigma_R^2}(r-\mu_R)^2 \right] \sum_{j=1}^{n} \frac{1}{\sqrt{2\pi}\varepsilon} \exp\left[-\frac{1}{2\varepsilon^2}(r-z-s_j)^2 \right] dr \qquad (6.13)$$

Exchange the order of the summation and integration in Equation 6.13, then simplify the result to obtain:

$$f_Z(z) = \frac{1}{n}\sum_{j=1}^{n} \frac{1}{\sqrt{2\pi}\sqrt{\varepsilon^2+\sigma_R^2}} \exp\left[-\frac{1}{2(\varepsilon^2+\sigma_R^2)}(z-(\mu_R-s_j))^2 \right], \quad -\infty < z < \infty$$

$$(6.14)$$

149

Equation 6.14 represents the KDE of the data $\mu_R - s_j$, $j = 1...n$, with smoothing

constant, $\varepsilon^2 + \sigma_R^2$. The CDF of Z is the integral of Equation 6.14:

$$F_Z(z) = \frac{1}{n} \sum_{j=1}^{n} \Phi\left[\frac{z - (\mu_R - s_j)}{\sqrt{\varepsilon^2 + \sigma_R^2}}\right], \quad -\infty < z < \infty \qquad (6.15)$$

where $\Phi(.)$ is the CDF of a standard normal random variable. The reliability of the

system is obtained by setting $z = 0$ in Equation 6.15 and observing

that, $reliability = 1 - F_Z(0)$ or $reliability = 1 - p_f$, then

$$reliability = 1 - \frac{1}{n} \sum_{j=1}^{n} \Phi\left[\frac{-(\mu_R - s_j)}{\sqrt{\varepsilon^2 + \sigma_R^2}}\right] \qquad (6.16)$$

Finally, the probability of failure of the system is given by:

$$p_f = \frac{1}{n} \sum_{j=1}^{n} \Phi\left[\frac{-(\mu_R - s_j)}{\sqrt{\varepsilon^2 + \sigma_R^2}}\right] \qquad (6.17)$$

The formulation presented above is now applied to the results obtained in Chapter 5 and a

comparison of the different methodologies used to quantify epistemic uncertainty relative

to this metric are shown.

6.3 Implementation and Results

In this section, the implementation of the formulation shown above, to calculate the

probability of failure given a collection of simulations obtained using the Bayes network

and the different methods to quantify interval data is presented. The first step in

calculating this probability of failure is to establish an acceptable threshold for this

system. Normally, this threshold will be specified a priori in a standards' manual, by historical testing, by a panel of experts and/or a combination of all of these. For this research, a suitable value is suggested based on some experimental evidence and expert opinion. It is then accepted that this threshold in itself also contains uncertainty which will not be treated in this research. The threshold is assumed to follow a Gaussian distribution with a mean of 247 g and a standard deviation of 34. The PDF of the failure threshold is shown in Figure 6.2.
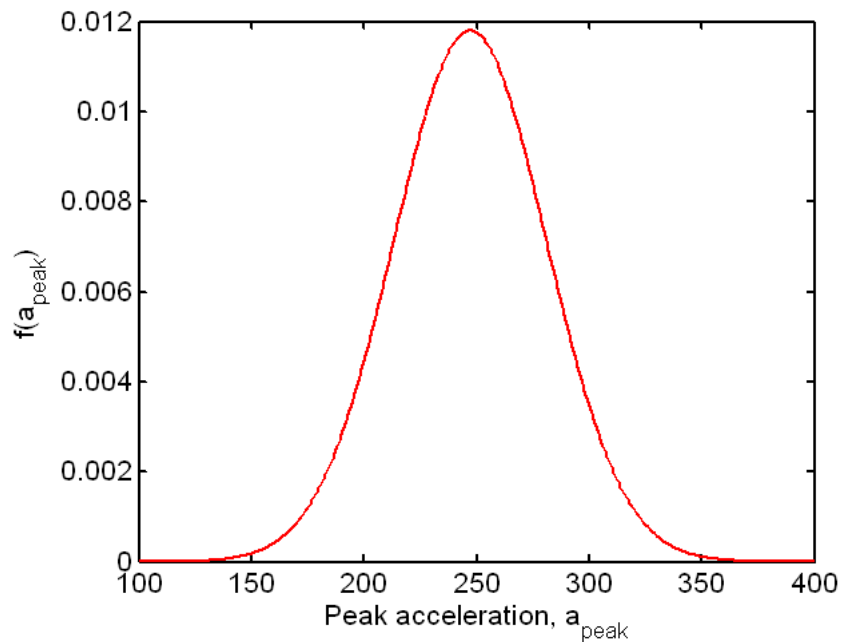
Figure 6.2. PDF of the system's peak acceleration threshold

Once this is established, the next step is to calculate the probability of failure for each of the simulated peak acceleration responses at the system level given that both epistemic and aleatoric uncertainty are present in the analysis. For this, the results at the system

level presented in Chapter 5 are used. The results for the percentile matching technique (presented in Section 5.4.2) are used to illustrate the methodology to calculate the probability of failure. The generated PDFs at the system level (first shown in Figure 5.16) and the system threshold PDF are shown in Figure 6.3. The overlap between the simulations and the threshold is what defines the probability of failure. This will be calculated using the formulation described in Section 6.2.
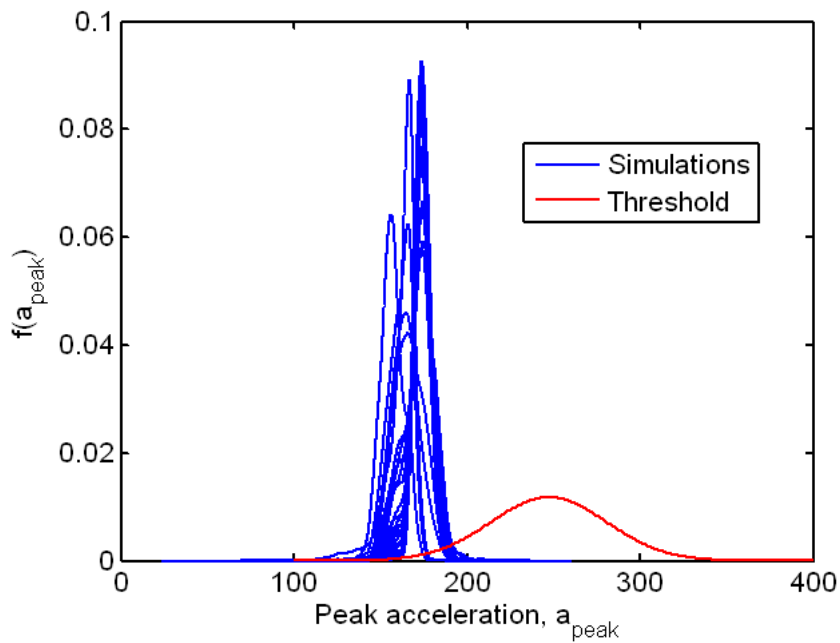
Figure 6.3. KDEs of simulated peak acceleration responses (blue), and PDF of threshold (red)

To simplify visualization of the steps involved in the calculation of the probability of failure, one of the realizations of the system response PDF (shown as blue curves in

Figure 6.3) is plotted along with the PDF of the threshold and they are shown in Figure 6.4.
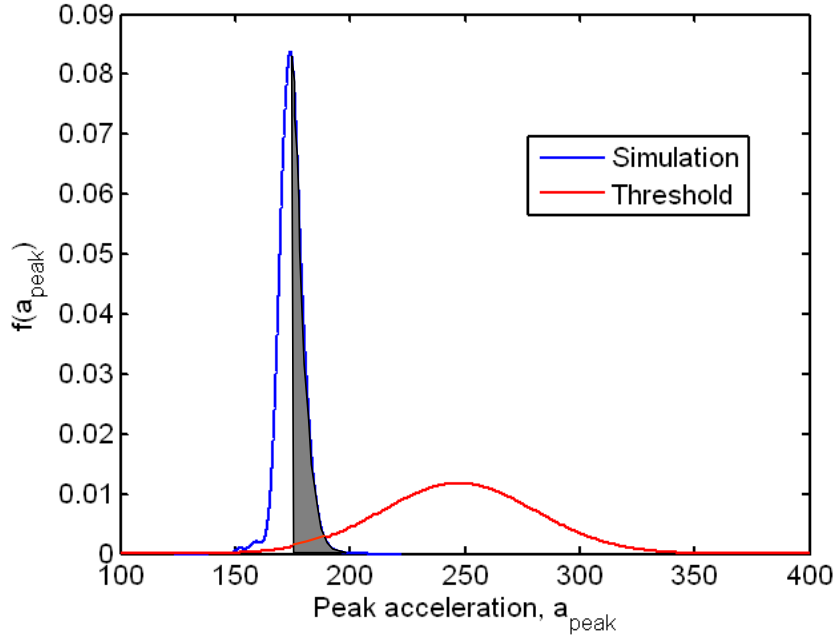
Figure 6.4. One realization of the PDF of absolute acceleration of system level response (blue) and the PDF of absolute acceleration threshold (red)

In this figure, the grey shaded area defines the probability of failure for this particular realization of the system level response. If the random variable, R (which in this case is the threshold of the system), falls in the interval $[r, r+dr]$ ($r = 175$ in the graph) and the random variable S (which comes from the Bayes network analysis) falls anywhere in the interval $[r,\infty]$ then failure occurs. The probability of this event is $f_R(r)(1 - F_S(r - z))dr$. To obtain the overall probability of failure, we integrate this event probability of failure over all r. This is what is done in Equation 6.8 with $z = 0$.

Using the formulation described in the previous section, the probability of failure for the case shown in Figure 6.4 is evaluated as 0.0170. The above procedure is repeated for each simulated PDF of S and the system threshold (shown as blue and red curves, respectively in Figure 6.3), and the corresponding failure probabilities are shown in Figure 6.5. Again, these failure probabilities were obtained at the system level using the percentile matching method. Note that these probabilities of failure are much higher than one would hope to attain for a high consequence system (where probabilities of failure are typically in the 1e-6 range). The reason for this is that the threshold was set arbitrarily low in order to have some overlap in the system response and threshold PDFs which facilitates the visualization of the results. The red colored triangle in Figure 6.5 represents the result of the case shown in Figure 6.4. As can be seen in the figure, most of the results are between 0.01 and 0.02 with one as low as 0.0045.
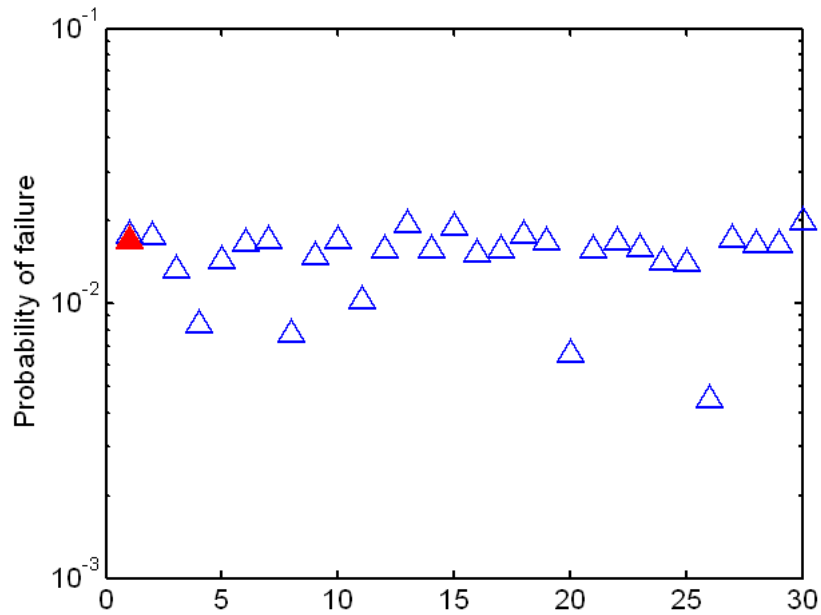
Figure 6.5. Collection of probabilities of failure from system level simulations using percentile matching. The red triangle represents the case shown in Figure 6.4.

Similarly for the system level response PDFs obtained using the other 5 methods for handling interval data presented in Chapter 5, the probability of failure is calculated. The values of the probability of failure will be greatly influenced by the choice of prior distribution of modulus of elasticity. As a matter of fact, the results are influenced by the method used to generate the prior distributions. In Figure 6.6, the resulting probabilities of failure obtained using the different methodologies presented in Chapter 5 to model interval data are presented. The legend in Figure 6.6 relates to the techniques presented in Chapter 5 as:

Methods for aggregating interval data

- 2 moment – Moment matching and bounding method

- Percentile – Percentile matching and 1st moment bounding

Methods for individual treatment of interval data

- Equal weight – Equal weighing of individual expert intervals

- Different weight – Different weights applied to each individual expert intervals

- Unif. Intervals – Each expert interval is modeled with a uniform distribution

- One uniform – One uniform distribution with bounds defined by the min/max of the expert-specified information.
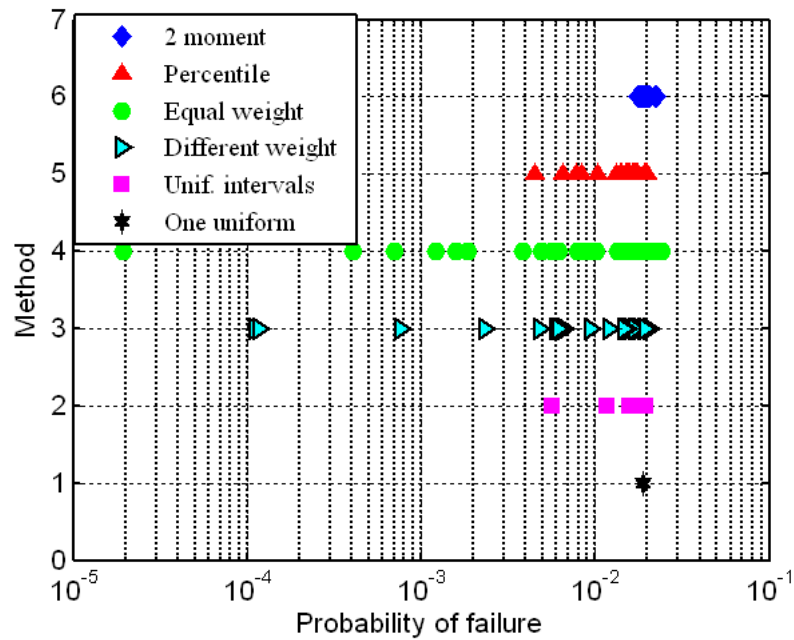
Figure 6.6. Comparison of probability of failure results for all methods of quantifying interval data. (Note log scale on the abscissa)

Some immediate conclusions can be drawn from Figure 6.6. For the case of one uniform distribution between the minimum and maximum of the bounds given by the experts, there is only one probability of failure calculated. This is analogous to the case where only aleatoric uncertainty is present (as described in Chapter 4). When compared to the other results, it is obvious what effect interval data has on the probability of failure. It expands the range of values that the probability of failure can occupy. It is this range of values that one should seek to minimize in order to increase the confidence that one has on a system level reliability prediction. This issue, as well as its ramifications connected to decision making, will be considered in Chapter 7

Another observation from Figure 6.6 concerns the spread of possible values for the probability of failure relative to the method used to create prior distributions. The smallest spread occurs when using the method of moment matching. This could be due to the priors' 1^{st} and 2^{nd} moments being constrained to match a given set and for the 3^{rd} and 4^{th} moments falling within some bounds obtained from the expert intervals. This may have the effect of constraining the form of priors that can be generated as well as where the majority of their density is located.

The widest range of values for the probability of failure is obtained when all the experts' intervals are equally weighted and plausible values of the parameters of the bounded Johnson distribution are obtained. The requirements placed on these distributions are that they lie within each of the expert-specified intervals and that the skewness and kurtosis, as defined in Equation 5.5 and 5.6, fall in the region identifying the random source as a bounded Johnson system as shown in Figure 5.1 and denoted S_B. This allows for a wider range of possible prior distributions to be generated and for some

157

of them to overlap, more or less, with the given threshold. These conditions will yield wider bounds on the probability of failure.

When different weights are assigned to each expert, the probability of failure has a range that is narrower than the one obtained for an equal weighing. This is not completely unexpected since the priors being used for the simulations tend to be biased toward one of the expert-specified intervals (as noted in Chapter 5, the relative weight of each expert is defined by the number of prior distributions from each interval; the higher the weight, the more priors defined from a particular interval). Again, for this particular set of weights for each experts' interval, the range of probability of failure is narrower than the equal weight case (for this particular problem). It is expected that a different set of weights will produce a different set of results, thus the above observation is not a generalization.

Finally, it is observed from Figure 6.6 that the results obtained from both the percentile matching and uniformly distributed prior analyses yield similar ranges of values for the probability of failure. The reason is as follows: if each of the values in an interval has equal probability of occurrence, then the range of the probabilities of failure is almost the same as if one constructs a bounded Johnson prior constrained to have its 10^{th} and 90^{th} percentile values within bounds of empirical CDFs arising from expert-specified bounds.

6.4 Conclusion

This chapter presented a reliability-based method to implement the QMU methodology. The metric of system sufficiency is the probability of failure. The

probability of failure compares random response levels to a given threshold of performance and includes sources of uncertainty. Both epistemic and aleatoric uncertainties are included in the system level response and aleatoric uncertainty is contained in the threshold. It is assumed, for the purpose of this research, that the threshold can be a random process since the actual system and its operating environment are random in nature.

The results shown in this chapter also highlight the effect of method choice in handling interval data; this is similar to model form uncertainty. This is reflected by the variation in results that occurs because of the use of different methodologies to model the interval variable, foam modulus of elasticity. As shown in the results, each technique yields a spread in the probability of failure which is attributed to the presence of interval data at the parametric level. The next chapter presents a methodology to assess the effects of two sources of epistemic uncertainty (interval data and method choice) within a decision making context.

CHAPTER VII

RESOURCE ALLOCATION USING QMU

7.1 Introduction

This chapter focuses on developing a methodology to allocate the resources needed to increase the confidence in the system level prediction. The method presented in the following sections leverages the work presented in Chapters 3 through 6. Based on a requirement on the amount of uncertainty present in a metric such as the probability of failure (presented in Chapter 6), an assessment could be made regarding whether or not a system is "certified" relative to that requirement in the presence of both epistemic and aleatoric uncertainty. This is what QMU seeks to address. If the answer is no, an important question to address is how to reach this requirement. This could be in terms of the resources needed to achieve a required level of confidence. Confidence in this dissertation is related to the amount of uncertainty present in the probability of failure of the system relative to a given threshold. Resources could be in terms of more experiments at a certain level, additional model simulations or model refinement and a reduction in the uncertainty of a metric of interest (such as the probability of failure) is equated to an increase in the confidence in the system level model prediction. Determining which resources are relevant could be done by means of a Phenomenology Identification and Ranking Table (PIRT), which connects the application requirements to some relevant phenomenon (Pilch et al., 2001) and Trucano et al., 2002).

To address the resource allocation question, a solution could be given by solving an optimization problem in which the design variables are functions of the resources to be allocated. In this study, analysis will be done numerically by introducing perturbations of one or more parameter(s) or one or more nodes in the Bayes network while keeping the others at their nominal values and calculating the metric of interest (i.e. probability of failure) for each combination.

7.2 Methodology

7.2.1 General approach

The general framework to address the resource allocation problem is now presented. This framework needs to include sources of aleatoric and epistemic uncertainty and will be cast as an optimization problem. In mathematical terms, this can be written as:

$$\min_{\Theta_i} \kappa_j$$

$$i = 1...n\_design\_variables$$
$$j = 1...n\_objectives$$

where
$$\kappa_j = f_j(\Psi)$$

$$\Psi = g(\Theta_i)$$

subject to

$$\text{cost} \leq \text{total budget}$$

$$\text{cost} = \text{cost of addl. resources}$$
$$= \sum_{i=1}^{nd} \alpha_i \Theta_i$$

and
$$\Theta_i^{lower} \leq \Theta_i$$

(7.1)

The various terms in Equation 7.1 are now described.

Θ_i is the i^{th} design variable which is constrained to be greater than Θ_i^{lower} (this is needed to avoid nonsensical cases, such as zero or negative number of data points); *n_design_variables* is the number of design variables that can be modified and have an effect on the uncertainty in the system performance. The design variables, Θ_i define where resources can be allocated to impact the confidence on the system model. These design variables could include:

1. Increased overall budget for testing, where permitted, at the different levels that make up the hierarchical system level model. In the example problem used in this study, this will be additional testing of joints and/or foam at levels 1 and 2.

2. Additional refinement to the computational models to reduce the error terms associated with the difference between the model predictions and the available experimental data at a particular level. For example, a functional form that relates the change in κ_j at the system level as a function of the reduction in error due to model refinement is needed.

3. Alternate models to describe the relevant physics of the problem. Again, a functional form that relates the change in κ_j at the system level as a function of the candidate physics models and their contribution to the error at each level will be required.

In Equation 7.1, Ψ is the measure of the system performance which is relevant to the application space of the system. It is directly related to the system level prediction and contains the different sources of uncertainty including variability due to part-to-part

variations and uncertainty due to interval data. This quantity incorporates the design threshold of the system which is defined a priori. In the case where interval data is present, such as in the example problem for this study, Ψ becomes a vector value as was shown in Chapter 6.

In Equation 7.1, κ_j is referred to as the system assessment metric and is a function of Ψ. It relates to the confidence in the system model. This quantity can take on various forms such as the range of Ψ and is defined as:

$$\kappa_j = \max(\Psi) - \min(\Psi) \tag{7.2}$$

or the expected value of Ψ given as:

$$\kappa_j = \frac{1}{n_\Psi}\sum \Psi$$
$$n_\Psi = \text{number of elements in } \Psi \tag{7.3}$$

Other candidate metrics are the expected value of information (EVI) (Hubbard, 2007) and entropy.

The constraints for this problem are specified as a function of the sum of the cost of each design variable ($\alpha_i \Theta_i$) where each α_i represents the unit cost of each design variable. The total cost should be less than or equal to a given budget.

In this study, the emphasis will be placed on how much experimental data needs to be included in the original problem. This translates into allocating resources (i.e. money) to run the necessary experiments at levels 1 and 2. This approach is examined in the next section.

7.2.2 Problem-specific approach

Following the general approach presented in the previous section, the specific formulation to the example problem used in this study is shown below. Using Equations 7.1 through 7.3, the following optimization problem can be written for the specific problem used in this study:

$$\min_{\Theta_i} \{\kappa_1, \kappa_2\}$$

$$i = 1, 2$$
$$j = 1, 2$$

$$\Theta_1 = n_1^f, \quad \Theta_2 = n_2^f$$

where
$$\kappa_1 = \text{range}(\Psi)$$

$$\kappa_2 = \text{E}[\Psi]$$

$$\Psi = p_f$$

subject to

$$\text{cost} \leq \text{total budget}$$

$$\text{cost} = \text{cost of addl. resources}$$
$$= \alpha_1 \Theta_1 + \alpha_2 \Theta_2$$
and
$$\Theta_1^{lower} \leq \Theta_1, \quad \Theta_2^{lower} \leq \Theta_2 \qquad (7.4)$$

where n_1^f and n_2^f are the optimal number of samples that minimizes the objective functions, κ_1 and κ_2. Ψ is a function of the number of foam samples at level 1 and 2 and the corresponding probability of failure (p_f). As it turns out, it is also a function of the method used to quantify the uncertainty due to interval data. Finding Ψ when interval data is present was the subject of Chapter 5 and the formulation presented in that chapter will be used in here. κ_1 was initially mentioned in Chapter 6 and was formally defined in Equation 7.2 as the range of the probability of failure at the system level. As noted in

Chapter 6, the range in values of probability of failure is a reflection of having a model

parameters specified in terms of interval data. It is also related to the method used to treat

the interval data probabilistically. κ_2 is the mean of the probability of failures based on

multiple realizations of the epistemic variable, E. The decision to minimize

κ_1 and κ_2 simultaneously is proposed in this study as a reasonable approach to perform

risk-informed decision analysis, as interpreted in this dissertation. The constraints for this

problem are stated in terms of the total budget for adding resources (i.e. more

experimental data). The values for each of the terms in Equation 7.4 are summarized in

Table 7.1.

Table 7.1. Summary of parameters in Equation 7.4

Parameter	Value/Description
Θ_1	Number of Level 1 foam samples
Θ_2	Number of Level 2 foam samples
Θ_1^{lower}	Minimum Level 1 foam samples = 3
Θ_2^{lower}	Minimum Level 2 foam samples = 3
Ψ	Probability of failure
κ_1	Range of probability of failure
κ_2	Mean of probability of failure
Total Budget	$100,000
α_1	$2,500 (per L1 sample)
α_2	$4,500 (per L2 sample)

To address the resource allocation problem, a multi-objective optimization problem (described in Equation 7.4) is solved. It involves the cost of additional foam samples at level 1 and 2 and requires the minimization of both κ_1 and κ_2. For completeness, the following is a brief summary of the available techniques to solve this problem (from Rao (1996)). These methods fall under the general category of direct search methods and are applicable to this problem.

1. Random Search Method: This method generates trial solutions for the optimization model using random number generators for the decision variables. Random search method includes random jump method, random walk method and random walk method with direction exploitation. Random jump method generates huge number of data points for the decision variable assuming a uniform distribution for them and finds out the best solution by comparing the corresponding objective function values. Random walk method generates trial solution with sequential improvements which is governed by a scalar step length and a unit random vector. The random walk method with direct exploitation is an improved version of random walk method, in which, first the successful direction of generating trial solutions is found out and then maximum possible steps are taken along this successful direction.

2. Grid Search Method: This methodology involves setting up of grids in the decision space and evaluating the values of the objective function at each grid point. The point which corresponds to the best value of the objective function is considered to be the optimum solution. A major drawback of this methodology is that the number of grid points increases exponentially with the number of decision variables, which makes the method computationally costlier.

3. Univariate Method: This procedure involves generation of trial solutions for one decision variable at a time, keeping all the others fixed. Thus the best solution for a decision variable keeping others constant can be obtained. After completion of the process with all the decision variables, the algorithm is repeated till convergence.

The majority of the methods described above are implemented in Matlab's optimization toolbox. One complicating factor for implementing these is that the design variables (i.e. n_1^f and n_2^f) are integers and thus cannot be directly accommodated in the Matlab optimization functions since these operate on continuous variables. After some initial attempts to solve this in Matlab using some workarounds, a decision was made to use a grid search technique to explore the possible design space. Even though this is the most expensive way of solving this problem, the fact that there are only 2 design variables and the design space is relatively small, made this a viable alternative to addressing this particular example.

7.3 Implementation

The solution implemented for this objective is shown graphically in Figure 7.1.
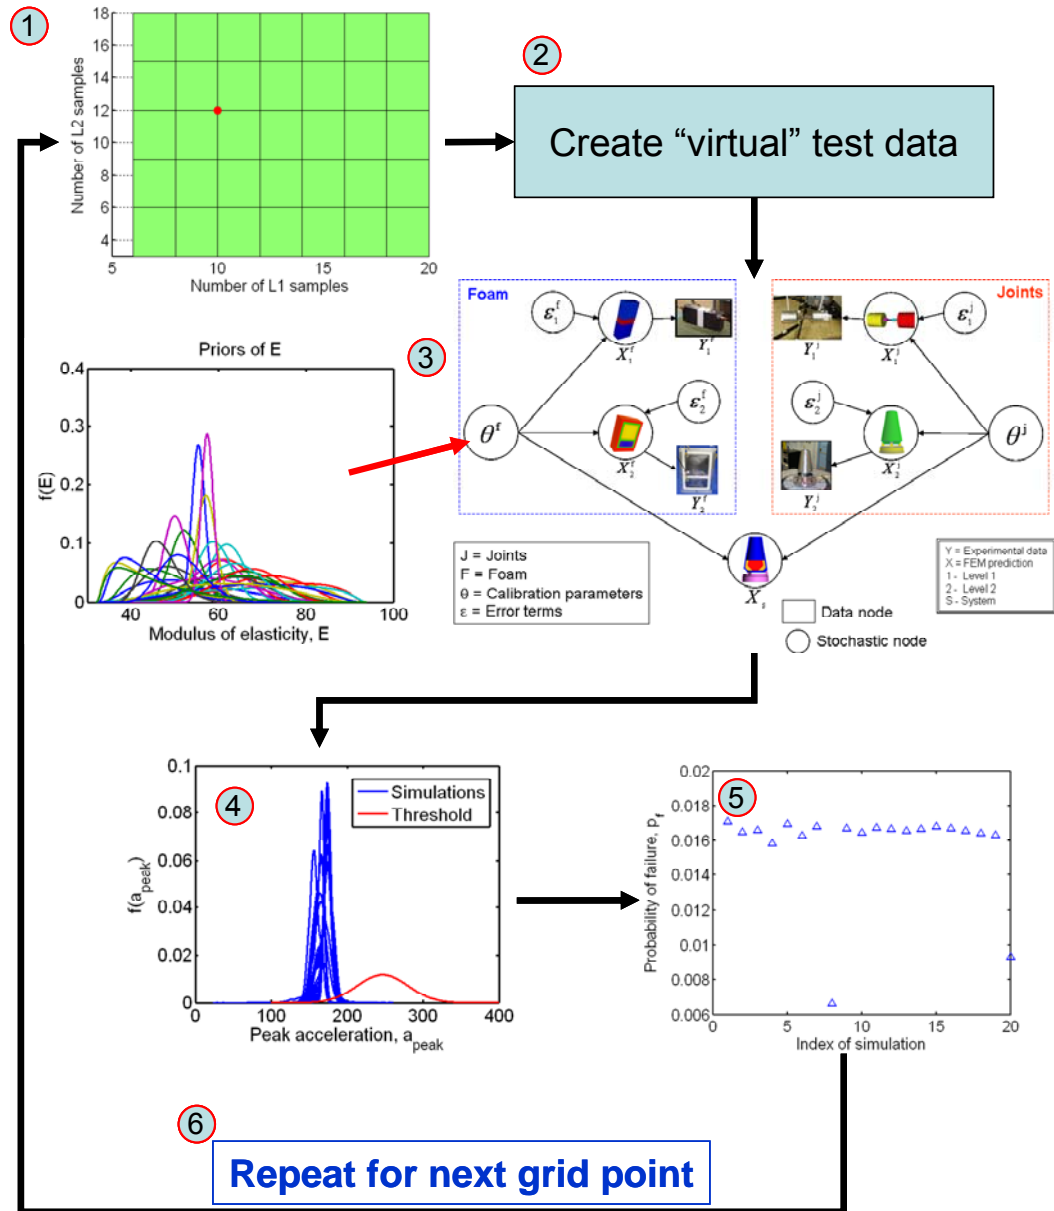
Figure 7.1. Schematic of grid search solution to resource allocation under uncertainty

The following steps detail the procedure to implement the resource allocation methodology shown in Figure 7.1. The step numbers shown below correspond to the numbers in the schematic in Figure 7.1

1. The first step is to define a reasonable space for the design variables n_1^f and n_2^f. This is shown in Figure 7.2 and it is usually defined by economics, availability of manufacturing and/or testing resources. For this research, it is assumed that all of these factors are considered when making an actual selection of the design space and the values used here are representative values chosen for illustration purposes only. Due to the large computational expense to calculate Ψ at all the combination of the points in the grid, the intervals in the grid were in increments of 2 and 3 for level 1 and 2 foam samples respectively. In addition and due to computational requirement in WinBUGS, the lower bound on the number of samples was set to 3.
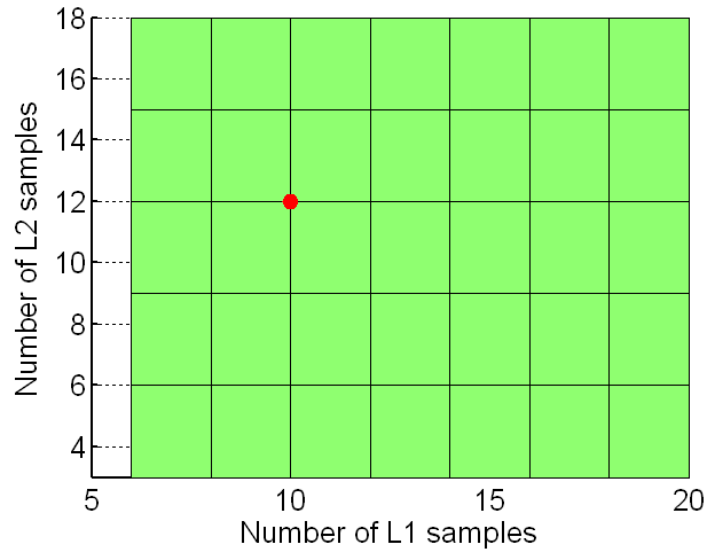
Figure 7.2. Grid of possible combinations of samples for level 1 and level 2 foam samples. Red dot shows one such combination.

2. An important issue to consider is how to obtain the data that is needed for updating the Bayes network developed in Chapter 4. Recall from Chapter 3 that only 6 and 3 data points are available from foam level 1 and 2 respectively for updating. Since data cannot be physically obtained instantaneously, one is faced with using "virtual" data to enable the necessary calculations. It is acknowledged that this comes with its own set of uncertainties which include the approach by which "virtual" data is generated. In this study, an assumption is made that the available level 1 and 2 foam data is representative of all possible scenarios and that it come from a uniform distribution. This uniform distribution has bounds defined by the minimum and maximum of the available data. The effect of this "virtual data" on the final decision is further investigated later in this chapter.

3. The next step is to update the Bayes network developed in Chapter 4 with the appropriate number of foam samples chosen from the grid shown in Figure 7.2. Since interval data is being considered, the process needs to be repeated for a certain number of prior distributions of modulus of elasticity which are generated by one (or all) of the methodologies described in Chapter 5. In this study and due to time constraints, 20 realizations of prior distributions of modulus of elasticity were used.

4. Once the Bayes network has been updated for all the priors using all the available data at all levels, a collection of posterior system level predictions are obtained and used with the given threshold to calculate the probability of failure of the system relative to peak acceleration. This process was described in Chapter 6 and defines the vector, Ψ in Equation 7.3.

5. The vector Ψ is shown in the schematic in Figure 7.1 as item #5. There are 20 values of the probability of failure which corresponds to each realization of the system level posterior distribution.

6. The process is repeated for all the combination of level 1 and 2 points in the grid and a collection of surfaces relating n_1^f, n_2^f and the probabilities of failures obtained after updating each of the priors of the modulus of elasticity can be plotted. In this case, 20 surfaces relating n_1^f, n_2^f and probability of failure are obtained.

The steps described above are used to populate the grid. Finding the optimum n_1^f and n_2^f by minimizing the objective functions described in Equations 7.2 and 7.3 are accomplished with the following procedure.

1. First, the values of α_1 and α_2 are used to calculate the cost of each of the combinations of n_1^f and n_2^f in the grid. Next the total budget constraint is imposed and only the combinations of n_1^f and n_2^f that satisfies the budgetary constraint are kept.

2. Using only the n_1^f and n_2^f that satisfy the budget constraint, their corresponding values of Ψ are used in Equations 7.2 and 7.3 to calculate κ_1 and κ_2 . These are the two objectives to minimize. To simultaneously minimize κ_1 and κ_2 , a weighted sum approach is taken were it is assumed that both objectives are equally weighted; thus, the minima is obtained by:

$$\min\left(\kappa_1 + \kappa_2\right) \tag{7.5}$$

3. With the result of Equation 7.5, the corresponding optimal n_1^f and n_2^f can be obtained and an associated cost can be computed.

With the procedure outline above, an estimate of the cost to minimize two objectives when interval data is present can be obtained. In this research, the effect of UQ method uncertainty when treating interval data was also investigated. To account for this, the process described above is repeated using four of the methods described in Chapter 5. These results are presented in Figure 7.3 through Figure 7.6 and some remarks regarding these results are made. The sequence of figures presented next consists of:

1. the collection of failure probabilities obtained from 20 realizations of the prior distributions of E for each method of treating interval data and at each grid point,

2. the value of κ_1 and κ_2 as defined in Equations 7.2 and 7.3,

3. a plot of κ_1 vs. κ_2 from which a point which minimizes both objectives as shown in Equation 7.5 can be obtained and

4. the corresponding optimal values of n_1^f and n_2^f and its associated cost.

In order to visualize κ_1 and κ_2 as a function of the design variables with a finer resolution than originally specified, and to avoid calculating the grid at each point, a Gaussian process model was used to represent each of the κ_1 and κ_2 surfaces as a function of n_1^f and n_2^f.

7.4 Results

7.4.1 Effect of epistemic UQ method uncertainty

This section shows the effect of the method used to model uncertainty on the overall system level uncertainty.

Figure 7.3 shows the results of the optimization when the moment matching method is used. From the collection of failure probabilities, the range and mean value of the probability of failure at each of the grid points can be calculated. The resulting optimal point that minimizes both objectives is shown on the lower right hand plot as 8 and 6 level 1 and 2 samples respectively with an associated cost of $47,000. This is the cost that minimizes both objectives simultaneously.
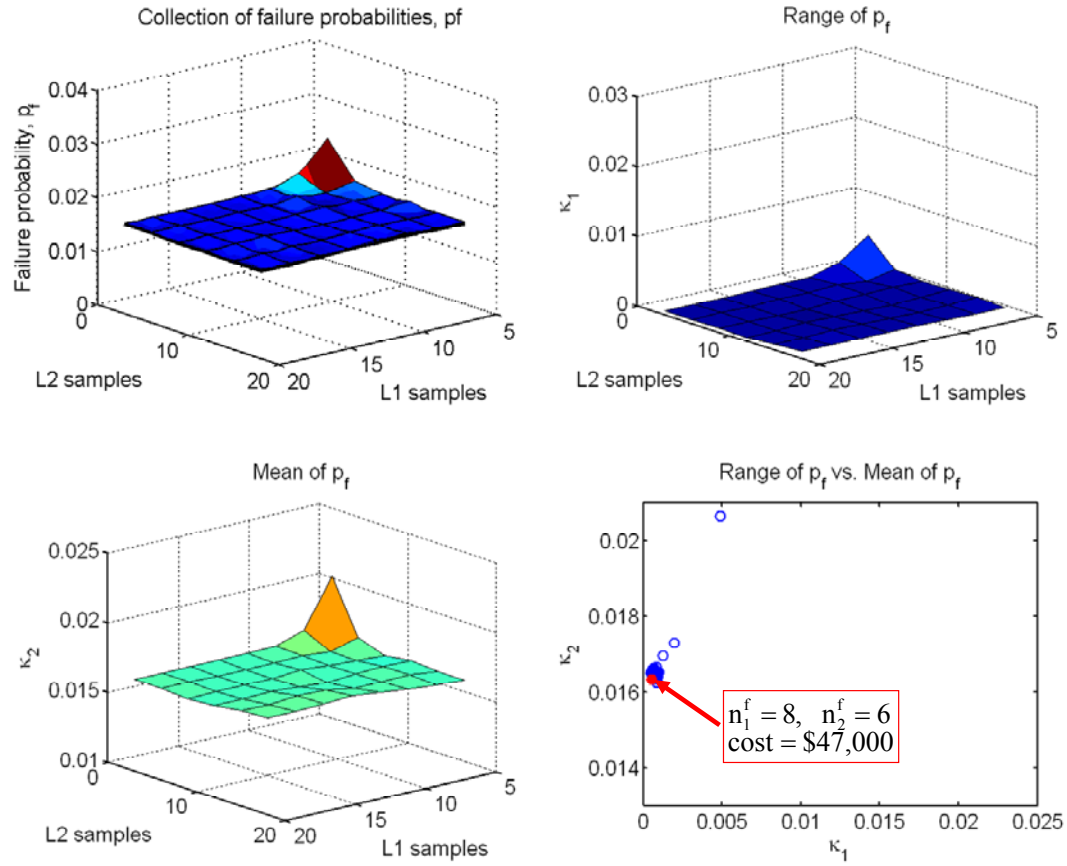
Figure 7.3. Results from implementing resource allocation methodology using moment matching method

Figure 7.4 shows the results of the optimization when the percentile matching method is used. Similarly to the above results, the optimal point that minimizes both objectives is shown on the lower right hand plot as 6 and 9 level 1 and 2 samples respectively with an associated cost of $55,500. This cost is higher than the one obtained with the moment matching method.
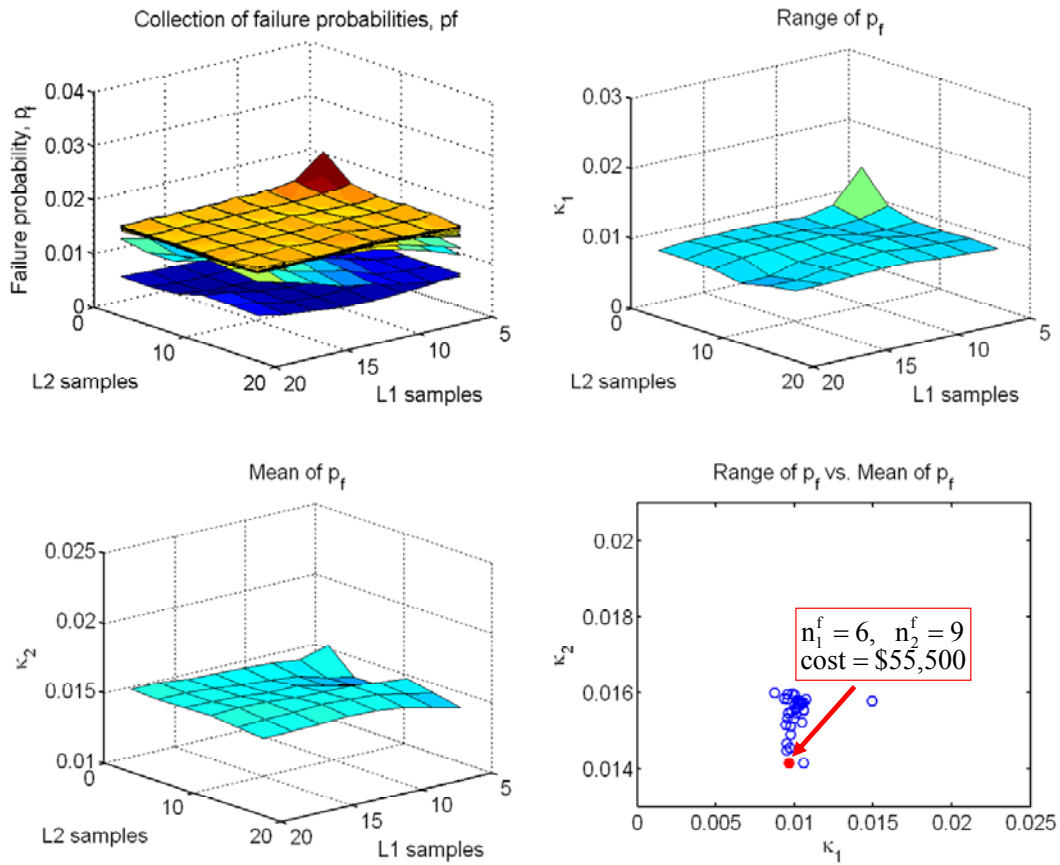
Figure 7.4. Results from implementing resource allocation methodology using percentile matching method

Figure 7.5 shows the results of the optimization when equal weighting of the experts is used. The optimal point that minimizes both objectives is shown on the lower right hand plot as 18 and 12 level 1 and 2 samples respectively with an associated cost of $99,500. This cost is very close to the total budget of $100,000 and higher than the previous methods.
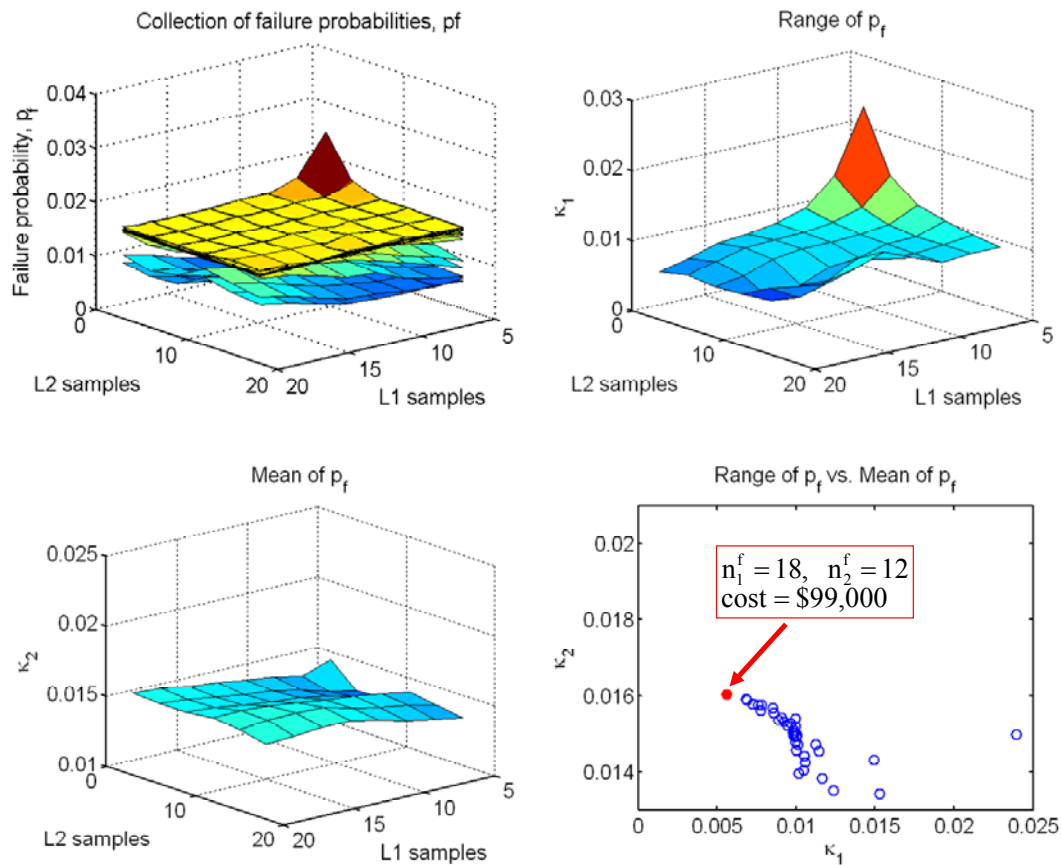
Figure 7.5. Results from implementing resource allocation methodology using equal weighing of experts

Finally, Figure 7.6 shows the results of the optimization when unequal weighting of experts is used. The optimal point that minimizes both objectives is shown on the lower right hand plot as 18 and 9 level 1 and 2 samples respectively with an associated cost of $85,500. Again, this cost is very close to the total budget of $100,000 but it is less than the previous method.
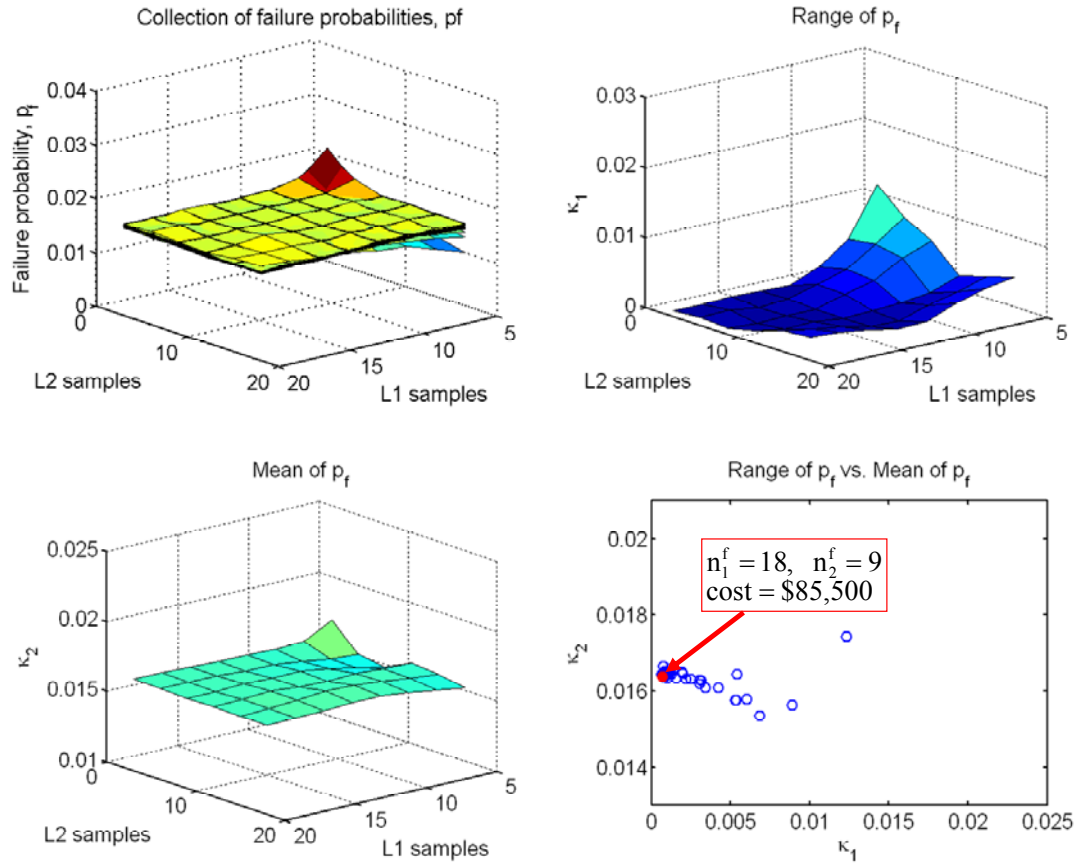
Figure 7.6. Results from implementing resource allocation methodology using different weighing of experts

The previous figures show the range of results that are obtained from using four of the methods to include interval data into the analysis. One encouraging feature in most of the plots is that some degree of convergence is achieved when enough test data is used to update the Bayes network. This convergence is shown as the flat region in each of the surfaces in the probability of failure plots. As expected the uncertainty in the results is high with limited data and starts to converge after a certain point. The optimal values for all four methods are summarized in Table 7.2.

Table 7.2. Comparison of optimal values using four methods to model interval data

Method	n_1^f	n_2^f	Cost ($)	κ_1	κ_2	$\kappa_1 + \kappa_2$
Moment matching	8	6	$47,000	0.000533	0.0163	0.0169
Percentile matching	6	9	$55,500	0.00967	0.0142	0.0238
Equal weighing	18	12	$99,000	0.0056	0.016	0.0216
Different weighing	18	9	$85,500	0.000674	0.0164	0.0171

It is obvious from Table 7.2 above that multiple solutions to the resource allocation problem are obtained and are dependent on the methodology used to quantify the epistemic uncertainty represented by interval data. For the analysis to be useful to a decision-maker, a single solution would be preferred versus a range of possible solutions. The method that provides the best solution to this problem, relative to the others, is based on the following criteria:

1. Minimizes the total cost

2. Minimizes the two objectives

3. Minimizes the range of failure probability (as this is related to the confidence of the system level prediction)

4. Gives the most conservative value of failure probability in a mean sense (i.e. the largest mean value at the optimal solution)

The above criteria are satisfied by the moment matching method and it is thus picked as the solution for this resource allocation problem.

7.4.2 Virtual experimental data

Based on the method selected above, attention is now focused on the issue of needing "virtual data" to enable the calculations needed to arrive at an optimal solution. As it was mentioned in the beginning of Section 7.3, "virtual experimental data" is created with the assumption that it comes from a uniformly distributed source and it is consistent with the limited available data for foam at level 1 and 2. To examine the effect of creating "virtual data" on the optimal solutions, several realizations of the "virtual data" were created and the optimization problem, using the moment matching method was solved for each realization. The results of these runs are shown in Figure 7.7 and summarized in Table 7.3.
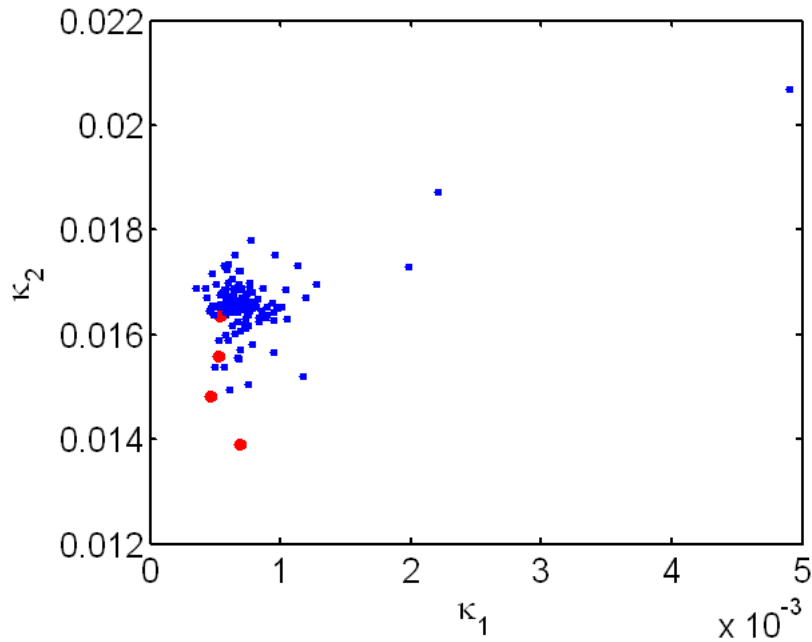
Figure 7.7. Comparison of results using multiple realizations of virtual experimental data

Table 7.3. Summary of optimal values using multiple realizations of virtual experimental data and moment matching method

Run #	n_1^f	n_2^f	Cost ($)	κ_1	κ_2	$\kappa_1 + \kappa_2$
1	8	6	$47,000	0.00053	0.0163	0.0169
2	8	3	$33,500	0.00068	0.0139	0.0146
3	14	3	$48,500	0.00053	0.0156	0.0161
4	12	6	$57,000	0.00046	0.0148	0.0153

The data shows the effect of four realizations of virtual data on the optimal solution of the problem. As it can be seen, there is no convergence to one optimal value but the range of values for the number of samples and the associated cost is not very

181

large. A solution could be stated as needing between \$33,500 and \$57,000 to obtain a range of failure probability between 0.00046 and 0.00068. This information accounts for the effect of multiple sources of uncertainty which include aleatoric uncertainty and the method by which epistemic uncertainty is modeled.

7.5 Conclusion

This chapter completes the implementation of RIDA which is a methodology to provide decision makers with information necessary to make informed decisions. In this study, those decisions involve allocating resources to increase the confidence in a system level prediction. Shown in Figure 7.1 is the schematic implementation of the resource allocation question to the demonstration problem used in this research. The range and the mean of the probability of failure are functions of the interval data and the number of experimental data points used for updating the Bayes network. They are also a function of the method used to model the interval data. This functional relationship was examined using a grid search to examine the design space of n_1^f and n_2^f which directly relate to the cost of building and testing additional foam samples at level 1 and 2. From these studies, a surface for each of the methods was obtained and both the range and the mean value of the probability of failure were minimized. From this, the corresponding optimal set of n_1^f and n_2^f and the cost associated with them can be obtained.

CHAPTER VIII

CONCLUSIONS AND FUTURE WORK

8.1 Summary of contributions

This dissertation described development of a methodology to enable risk-informed decision analysis (RIDA) that provides decision makers with information necessary to make informed decision regarding allocation of resources to increase the confidence in the prediction of a system level model of an engineering system. Central to the concept of RIDA is the quantification of margins and uncertainties (QMU) which in essence, seeks to quantify sources of both known and unknown conditions. These are commonly classified as aleatoric and epistemic uncertainties and the probabilistic treatment of the latter is a key component of this study.

To demonstrate a methodology that enables RIDA, a multi-level, 2-component problem developed at Sandia is used and was described in Chapter 3. The system level model in the Sandia problem was built in a hierarchical or building-block approach manner which builds complex system model using simpler component models. This approach takes advantage of data that is available at the component level. In this research, it is proposed that all available data be used to quantify the uncertainty at the system level. With this in mind, the schematic implementation of resource allocation to the demonstration problem described in Chapter 3 is shown in Chapter 7. It shows all the components that RIDA is comprised of:

1. Quantification of aleatoric and epistemic uncertainty and propagation to the system level.

2. Quantification of margins and uncertainty (QMU) at the system level.

3. Resource allocation under uncertainty.

The first component was addressed in Chapters 4 and 5 with the formulation of a Bayes network that incorporates all the available data at the different levels. Just to reiterate, there is no experimental data at the system level. In Chapter 5, the modulus of elasticity of foam was given in terms of interval data, a type of epistemic uncertainty. Modeling interval data in a probabilistic framework was the subject of Chapter 5. A flexible family of distributions was used as the basis for modeling the interval data. Several methods to determine the parameters of this family of distributions were examined. The effect of each method on the uncertainty at the system level was observed. This is a form of model form uncertainty since the true distribution of the interval data is unknown. This uncertainty as well as the uncertainty due to interval data is then reflected in the second component of RIDA.

The second component of RIDA, quantification of margins and uncertainties (QMU), was addressed in Chapter 6. This chapter proposes a reliability-type framework to address QMU. In this framework, sources of uncertainty are accounted for in both the system level response and in the specified requirement of the system. This requirement is typically given a priori. The formulation presented in Chapter 6 makes use of basic reliability concepts and uses kernel density estimators to calculate the probability of failure of the system. Since interval data is present and treated probabilistically, the probability of failure becomes a vector of values. Since model form uncertainty is also

present, the probability of failure becomes a matrix. In other words, the probability of failure becomes a function of both the interval data and the method in which interval data is modeled. This fact complicates the problem since now there is a range of possible values of probability of failure. For this information to be useful to decision-makers, a methodology that accounts for the uncertainty in the probability of failure was developed.

The third component, decision-making under uncertainty, was discussed in Chapter 7. As noted in the paragraph above, the probability of failure can have a range and a mean value due to interval data and to the method used to model it. The uncertainty present in both the range and mean of the probability of failure is influenced by the choice of method used to model the interval data and thus needs to be accounted for. To use this in a resource allocation context, this study proposed using an optimization approach where the design variables relate to those resources that have some impact on the uncertainty of the system level response. In this study, these resources were in terms of the cost to manufacture and perform additional experiments for the foam components at level 1 and 2. Analyses performed using the methodology developed in Chapters 4 through 6 eventually lead to a collection of optimal cost of adding foam data based on the uncertainties being included. It was obvious from the results, that in the presence of epistemic uncertainty, there is not a single solution to the problem but the information obtained with the analysis in Chapter 7 could help a decision-maker narrow down the possible alternatives when allocating or requesting resources.

8.2 Proposed future work

The current research addresses the fundamental question of how to incorporate different sources of data at different levels of complexity into a hierarchically built complex system level model. Both aleatoric and epistemic uncertainty was incorporated and propagated in a probabilistic manner up to the system level response of interest. Following the analysis done to address the four main objective of this research, a set of issues were identified and are listed below as possible research directions for future implementation:

1. In this research, the parameters of the Smallwood joint model were treated as statistically independent. Statistical correlations between the model parameters need to be quantified and included in the context of using a Bayes network.

2. When examining the example problem and the corresponding Bayes network, it is observed that some levels could be more relevant to the system level than others. This is because the geometry or the physics of interest are similar at a particular level relative to the system level. A question is: how to give different weights to different levels, considering which level is closer to the system level?

3. One of the cases examined to incorporate epistemic uncertainty assigned either equal or unequal weighing to the various experts. This weighing was done in an ad-hoc manner in this research. A more formal way to assign these weighing is required. This weighing scheme could be tied to the comparison measure (i.e. the probability of failure).

4. To further examine the effect of model form uncertainty (in the physical system model), the following question is raised: how to consider model uncertainty,

which comes from multiple models that contribute to some of the nodes? For example, the Smallwood model was considered in this research but others are available. How to assess the effect of this type of uncertainty?

5. Currently we only considered a single system output. What if there are multiple outputs of interest? One possibility is a system reliability approach for QMU.

6. Currently the optimization for resource allocation in Chapter 7 is done using a brute force grid search. There is a need for a more elegant algorithm that incorporates sensitivities of system output and uncertainty to various sources of uncertainty/error. Non-gradient based optimization methods could be good candidates for this.

7. A related issue stemming from item 6 above is the overall question of how to do resource allocation in the case of multiple outputs.

8. There is a need for efficient computation of the Bayes network. Particularly in the case of uncertainty quantification were the Bayes network might need to be run many times, it will be advantageous to make this computation as efficient as possible especially for a more complex system. Several ideas come to mind but one involves parallel processing of the nodes in the Bayes network. If each node could run on a separate processor, then one could take advantage of parallel computers to run this more efficiently. Another idea would involve the actual algorithm used for solving the Bayes network. Currently, MCMC-based solutions for the Bayes network are implemented and this requires convergence of the network which takes many iterations. Alternative algorithms that provided faster integration are desired.

REFERENCES

Alvin, K.F., Oberkampf, W.L., Rutherford, B.M. and Diegert, K.V., 2000, "Methodology for Characterizing Modeling and Discretization Uncertainties in Computational Simulation", SAND2000-0515, Sandia National Laboratories, Albuquerque, NM.

Ang, A. H-S. and Tang, W.H., 1975, *Probability Concepts in Engineering Planning and Design, Vol.I - Basic Principles*, John Wiley and Sons, New York, NY.

Aloise, G., 2007, *Nuclear Weapons: Annual Assessment of the Safety, Performance, and Reliability of the Nation's Stockpile*, GAO-07-243R Annual Assessment, U.S. Government Accountability Office, Washington, D.C.

Apostolakis, G., 1990, "The Concept of Probability in Safety Assessments of Technological System", Science, Vol. 250, No. 4986, pp. 1359-1364.

AIAA, 1998, Guide for the Verification and Validation of Computational Fluid Dynamic Simulations, American Institute of Aeronautics and Astronautics, AIAA-G-077-1998, Reston, VA.

Barlow, R. E., 1988, "Using Influence Diagrams", *Accelerated life testing and experts' opinions in Reliability*, pp.145–150.

Bichon, B.J., Eldred, M.S., Swiler, L.P., Mahadevan, S., McFarland, J.M., 2008, "Efficient Global Reliability Analysis for Nonlinear Implicit Performance Functions", to appear in AIAA …

Bobbio, A., Portinale, L., Minichino, M. and Ciancamerla, E., 2001, "Improving the Analysis of Dependable Systems by Mapping Fault Trees into Bayesian Networks", Reliability Engineering and System Safety, vol. 71(3), pp.249–260.

Cobb, B.R. and Shenoy, P.P., 2006, "On The Plausibility Transformation Method For Translating Belief Function Models To Probability Models", International Journal of Approximate Reasoning, Vol. 41, pp. 314–330.

Cox, D.R. and D. Oakes, 1984, *Analysis of Survival Data*, Chapman & Hall, London.

CUBIT, 2008, The CUBIT Geometry and Mesh Generation Toolkit, Sandia National Laboratories, http://cubit.sandia.gov/documentation.html.

DeBrota, David J., Swain, James J., Roberts, Stephen D., Venkataraman, Sekhar., 1998, "Input modeling with the Johnson System of distributions," *Proceedings of the 1988 Winter Simulation Conference.*

Diegert, K., Klenke, S., Novotny, G., Paulsen, R., Pilch, M. and Trucano, T., 2007, *Toward a More Rigorous Application of Margins and Uncertainties within the Nuclear Weapons Life Cycle – A Sandia Perspective*, SAND2007-6219, Sandia National Laboratories, Albuquerque, New Mexico.

Efron, B., and Tibshirani, R.J., 1998, *An Introduction to the Bootstrap*, Chapman & Hall/CRC Press, Boca Raton, FL.

Ferson, S., Kreinovich, V., Hajagos, J., Oberkampf, W., and Ginzburg, L., 2007, *Experimental Uncertainty Estimation and Statistics for Data Having Interval Uncertainty*, SAND2007-0939, Sandia National Laboratories, Albuquerque, NM.

Ferson, S., R.B. Nelsen, J. Hajagos, D.J. Berleant, J. Zhang, W.T. Tucker, L.R. Ginzburg and W.L. Oberkampf. 2004. *Dependence in Probabilistic Modeling, Dempster-Shafer Theory, and Probability Bounds Analysis*. SAND2004-3072, Sandia National Laboratories, Albuquerque, New Mexico.

Foster, J., Agnew, H., Gold, S., Guidice, S. and Schlesinger, J., 1999, "FY 1999 Report to Congress of the Panel to Assess the Reliability, Safety and Security of the Unites States Nuclear Stockpile"

Gelman, A. J. B. Carlin, H. S. Stern and D. B. Rubin, 2004, *Bayesian Data Analysis*, 2nd edition, Chapman and Hall/CRC, Boca Raton.

Gibson, L.J. and Ashby, M.G., 1999, Cellular Solids: Structure and Properties, 2nd edition, Cambridge University Press.

Gilks, G. R., Richardson, S., and Spiegelhalter, D. J., *Markov Chain Monte Carlo in Practice*, Interdisciplinary Statistics, Chapman & Hall/CRC, London, 1996.

Gill, J., 2002, *Bayesian Methods – A Social and Behavioral Sciences Approach*, Chapman & Hall/CRC Press, Boca Raton, FL.

Gioia, F., and C.N. Lauro. 2005. Basic statistical methods for interval data. *Statistica Applicata* [*Italian Journal of Applied Statistics*] 17(1): 75-104.

Gran, B. A. and Helminen, A., 2001, "A Bayesian Belief Network for Reliability Assessment", Safecomp 2001, vol.2187, pp.35–45

Haldar, A. and Mahadevan, 2000, S., *Probability, Reliability and Statistical Methods in Engineering Design*, John Wiley and Sons, New York.

Helton, J.C., Johnson, J.D., Oberkampf, W.L. and Storlie, C.B., 2006, "A Sampling-Based Computational Strategy for the Representation of Epistemic Uncertainty in Model Predictions with Evidence Theory", SAND2006-5557, Sandia National Laboratories, Albuquerque, NM.

Eardley, D., 2005, *Quantification of Margins and Uncertainties (QMU)*, JSR-04-330, JASON, The MITRE Corporation, McLean, Virginia.

Jensen, Finn V., 2001, *Bayesian Networks and decision graphs*, Springer-Verlag, New York.

Jiang, X. and Mahadevan, S., 2007, "Bayesian risk-based decision method for model validation under uncertainty", Reliability Engineering and System Safety, Vol. 92. pp. 707-718.

Johnson, N.L., 1949, "Systems of frequency curves generated by methods of translation", Biometrika, 36:149-176.

Hubbard, D., 2007, *How to Measure Anything: Finding the Value of Intangibles in Business*, John Wiley & Sons, New York, NY.

Kennedy, M.C. and O'Hagan, A., 2000, "Predicting the output from a complex computer code when fast approximation are available", Biometrika, Vol. 87, 1, pp. 1-13.

Kreinovich, V., L. Longpré, S. Ferson, and L. Ginzburg. 2004. *Computing Higher Central Moments for Interval Data.* University of Texas at El Paso, Department of Computer Science, Technical Report UTEP-CS-03-14b.

Landes, R.D., Loutzenhiser, P.G. and Vardeman, S.B., 2006, "Hierarchical Bayesian Statistical Analysis for a Calibration Experiment", IEEE Transactions on Instrumentation and Measurement, Vol. 55, No. 6, pp. 2165-2171.

McDonald, M., Zaman, K., Mahadevan, S., 2009a, "Representation and First-Order Approximations for Propagation of Aleatory and Distribution Parameter Uncertainty", In Proceedings of the 50th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, Palm Springs, California, Paper number AIAA-2009-2250.

McDonald, M., Zaman, K., Rangavajhala, S., Mahadevan, S., 2009b, "A probabilistic approach for representation of interval uncertainty", under review, Reliability Engineering and Systems Safety.

McDonald, M. and S., Mahadevan, S., 2009c, "Uncertainty Quantification and Propagation for Multidisciplinary System Analysis", In Proceedings of the 50th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, Palm Springs, California.

Mahadevan, S. and Rebba, R., 2005, "Validation of reliability computational models using Bayes networks", Reliability Engineering and System Safety, Vol. 87. pp. 223-232.

Mardia, K. and Marshall, R., 1984, "Maximum likelihood estimation of models for residual covariance in spatial regression", Biometrika, Vol. 71, pp. 135–146.

Martin, J. and Simpson, T., 2005, "Use of kriging models to approximate deterministic computer models", AIAA Journal, Vol. 43(4), pp.853–863.

McFarland, J.M., 2008, *Uncertainty analysis for computer simulations through validation and calibration*, Vanderbilt University, Nashville, TN.

Mourelatos, Z.P. and Zhou, J., 2006, "A Design Optimization Method Using Evidence Theory", Journal of Mechanical Design, Vol. 128, pp. 901-908.

Oberon Microsystems, Inc., 2006, *Component Pascal Language Report*, Switzerland.

Oberkampf, W.L. and Trucano, T.G., 2000, "Validation Methodology in Computational Fluid Dynamics", SAND 2000-1656C, Sandia National Laboratories, Albuquerque, NM.

Oberkampf, W.L., DeLand, S.M., Rutherford, B.M., Diegert, K.V., and Alvin, K.F., 2000a, "Estimation of total uncertainty in modeling and simulation", SAND 2000-0824, Sandia National Laboratories, Albuquerque, NM.

Osegueda, R., V. Kreinovich, L. Potluri, R. Aló. 2002. Non-destructive testing of aerospace structures: granularity and data mining approach. Pages 685-689 in *Proceedings of FUZZ-IEEE 2002*, Vol. 1, Honolulu, Hawaii.

Paez, T.L., 2009, personal notes and communications.

Parry, G.W., and Winter, P.W., 1981, "Characterization and Evaluation of Uncertainty in Probabilistic Risk Analysis", Nuclear Safety, Vol. 22, No. 1, pp. 28-42.

Pepin, J.E., Rutherford, A.C. and Hemez, F.M, 2008, "Defining a Practical QMU Metric", AIAA 2008-1717, Proceedings of the 49th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference, Schaumburg, IL.

Perkins, R.H., 2000, "Assuring the Safeth and Reliability of America's Nuclear Weapons: The Annual Certification Process", National Defense University National War College, Washington, DC.

Pilch, M., Trucano, T.G., and Helton, J.C., 2006a, "Ideas Underlying Quantification of Margins and Uncertainties (QMU): A White Paper", SAND2006-5001, Sandia National Laboratories, Albuquerque, NM.

Pilch, M., Trucano, T.G., Moya, J.L., Froehlich, G., Hodges, A. and Peercy, D., 2001, "Guidelines for Sandia ASCI Verification and Validation Plans – Content and Format: Version 2.0", SAND2000-3101, Sandia National Laboratories, Albuquerque, NM.

Rao, S.S., 1996, *Engineering Optimization – Theory and Practice*, 3rd edition, John Wiley and Sons, New York, NY.

Rasmussen, C., 1996, *Evaluation of Gaussian processes and other methods for non-linear regression*, PhD thesis, University of Toronto.

Rebba, R., 2005, *Model Validation And Design Under Uncertainty*, Vanderbilt University, Nashville, TN.

Rebba, R. and Mahadevan, S., 2006a, "Inclusion of Model Errors in Reliability-Based Optimization", Journal of Mechanical Design, Vol. 128, pp. 936-944,

Rebba, R., and Mahadevan, S., 2006b, "Model Predictive Capability Assessment Under Uncertainty", AIAA Journal, Vol. 44, No. 10, pp. 2376-2384.

Rebba, R., Mahadevan, S. and Huang, S., 2006c, "Validation and error estimation of computation model", Reliability Engineering and System Safety, Vol. 91. pp. 1390-1397.

Red-Horse, J.R. and Benjamin, A.S., 2004, "A Probabilistic Approach to Uncertainty Quantification with Limited Information", Reliability Engineering and System Safety, *Alternative Representation of Epistemic Uncertainty*, J.C. Helton and W.L. Oberkampf, guest editors, Vol. 85. Nos. 1-3, pp. 183-190.

Reese, G., Bhardwaj, M., Segalman, D., Alvin, K., Driessen, B., Pierson, K., Walsh, T., (1999), "Salinas – Users Notes," Sandia Report 99-2801, Sandia National Laboratories, Albuquerque, NM.

Reliability Engineering and System Safety (RESS), 2004, *Alternative Representation of Epistemic Uncertainty*, J.C. Helton and W.L. Oberkampf, guest editors, Vol. 85. Nos. 1-3, July-September.

Robinson, D., 2001, "A Hierarchical Bayes Approach to System Reliability Analysis", SAND2001-3513, Sandia National Laboratories, Albuquerque, NM

Santner, T. J., Williams, B.J. and Noltz, W. I., *The Design and Analysis of Computer Experiments*, Springer-Verlag, New York, 2003.

Sentz, K and Ferson, S., 2002, "Combination of evidence in Dempster-Shafer theory, SAND 2002-0835, Sandia National Laboratories, Albuquerque, NM.

Sharp, D.H. and Wood-Schultz, M.M., 2003, "QMU and Nuclear Weapons Certification", Los Alamos Science 28:47-53.

Silverman, B. W., 1986, *Density Estimation for Statistics and Data Analysis*, Chapman and Hall, London.

Sindir, M. M., Barson, S. L., Chan, D. C., and Lin, W. H., "On the Development and Demonstration of a Code Validation Process for Industrial Applications," American Institute of Aeronautics and Astronautics, AIAA Paper No. 96- 2032, *27th AIAA Fluid Dynamics Conf.*, New Orleans, LA, 1996.

Smallwood, D., Gregory, D., Coleman, R., 2000, "Damping Investigations of a Simplified Frictional Shear Joint," Proceedings of the 71st Shock and Vibration Symposium, SAVIAC, The Shock and Vibration Information Analysis Center.

Spiegelhalter, D. J., Thomas, A., Best, N. G. and Lunn, D., 2003, *WinBUGS User Manual Version 1.4*. Cambridge, U.K.: MRC Biostatistics Unit, [Online], Available: http://www.mrc-bsu.cam.ac.uk/bugs

Soundappan, P., Nikolaidis, E., Haftka, R.T., Grandhi, R. and Canfield, R., 2004, "Comparison of evidence theory and Bayesian theory for uncertainty modeling", Reliability Engineering and System Safety, Vol. 85. pp. 295-311.

The Mathworks, Inc, 2009, "Matlab: On-line documentation", Natick, Massachusetts.

Trucano, T.G., Easterling, R.G., Dowding, K.J., Paez, T.L., Urbina, A., Romero, V.J., Rutherford, B.M. and Hills, R.G., 2001, "Description of Sandia Validation Metrics Project, SAND 2001-1339, Sandia National Laboratories, Albuquerque, NM.

Trucano, T. G., Pilch, M., and Oberkampf, W. L., 2002, "General Concepts for Experimental Validation of ASCI Code Applications," Sandia National Laboratories, SAND2002-0341, Albuquerque, NM.

Urbina, A., Paez, T.L., Gregory, D., Resor, B., Hinnerichs, T.D. and O'Gorman, C.C, 2006, "Validation of a Combined Non-Linear Joint and Viscoelastic Encapsulating Foam", Proceedings of the 2006 Society for Experimental Mechanics, St. Louis, MO.

U.S. Department of Energy, 2000, "Advanced Simulation and Computing (ASCI) Program Plan," 01-ASCI-Prog-01, Sandia National Laboratories, Albuquerque, New Mexico.

Venkatraman, S. and Wilson, J.R., 1987, "Modeling Univariate Populations with Johnson's Translation System-Description of the FITTR1 software." Research Memorandum 87-21, School of Industrial Engineering, Purdue University, West Lafayette, Indiana.

Williams, B., Higdon, D., Gattiker, J., Moore, L., McKay, M., and Keller-McNulty, S., 2006, "Combining Experimental Data and Computer Simulations, With an Application to Flyer Plate Experiments", Bayesian Analysis, Vol. 1, Number 4, pp. 765-792.

Wilson, A.G. and Huzurbazar, A.V., 2007, "Bayesian Networks for multilevel system reliability", Reliability Engineering and System Safety, Vol. 92. pp. 1413-1420.

Wu, J.S., Apostolakis, G.E. and Okrent, D., 1990, "Uncertainties in system analysis: probabilistic versus non-probabilistic theories, Reliability Engineering and System Safety, Vol. 30. pp. 163-181.

Xiang, G., S.A. Starks, V. Kreinovich and L. Longpré. 2006. New algorithms for statistical analysis of interval data. *Springer Lecture Notes in Computer Science* 3732: 189-196.

Zhang, R. and S. Mahadevan, 2003, "Bayesian Methodology for Reliability Model Acceptance," Reliability Engineering and System Safety, Vol. 80, 95-103.